

NCC 2-387

AMES GRANT

IN-61-CR

127027

P-105

# Research Summary

January 1989 - June 1990

(NASA-CR-190960) RESEARCH SUMMARY,  
JANUARY 1989 - JUNE 1990 Report, 1  
Jan. 1989 - 30 Jun. 1990 (Research  
Inst. for Advanced Computer  
Science) 105 p

N93-13694

Unclass

G3/61 0127027

RIACS



# Research Summary

January 1989 - June 1990

**Peter J. Denning, Director 1983-1990**



---

# Table of Contents

Page

<b>Introduction .....</b>	<b>1</b>
<b>Overview .....</b>	<b>3</b>
<b>The Science of Computing.....</b>	<b>9</b>
<b>Learning Systems.....</b>	<b>13</b>
The Sparse Distributed Memory Project .....	13
Objectives of the Project .....	15
Strategies for Achieving the Objectives .....	15
Accomplishments .....	16
<i>Architectural Studies</i>	
SDM Alternative Designs.....	16
Speech Recognition .....	17
Digital Prototype of SDM.....	18
Cerebellar Modeling .....	18
<i>Sensory Encoding</i>	
Fine Discrimination Based on Distinctive Features.....	19
Encoding Features for SDM Image Recognition.....	19
<i>Applications and Technology Transfer</i>	
Two-Dimensional Shape Recognition .....	20
Large-Scale Simulation on the CM-2 Connection Machine .....	21
Weather Prediction.....	21
Applications of High-Dimensional Associative Memories to Robot Control .....	22
The Bayesian Learning Project.....	24
Discoveries in Astronomy .....	25
<b>Networked Systems.....</b>	<b>27</b>
Overview.....	27
NASA Sponsors and Collaborators .....	30
Accomplishments .....	31
<i>Distributed Systems</i>	
Distributed Intelligent Agents.....	31
Network Analysis.....	31
CCSDS Activities .....	32
Space Station Freedom Testbed Activities.....	33
<i>Collaboration Technology</i>	
NASA Science Internet Project .....	33
Remote Coaching Facility.....	34

---

<i>Workstation Technology</i>	
Workstation Development.....	34
Workstation Workshop .....	35
Virtual Planetary Exploration Project.....	35
<i>Applied Pilot Projects</i>	
Telescience Coordination and Planning .....	36
Remote Coaching Facility.....	37
Global Environmental Change Project.....	37
User Support at Supercomputer Facilities.....	38
Unified Information System Project .....	38
Parallel Processing Database .....	38
<b>Parallel Systems</b> .....	41
Objectives.....	42
Accomplishments .....	45
<i>Parallel Linear Algebra Algorithms</i>	
Efficient Preconditioned Conjugate Gradient on Parallel Computers .....	45
Sparse Solution Algorithms for CFD Applications .....	46
Study of Structured and Partitioned Indefinite Hermitian Matrices and the Development of Preconditioners for such Matrices.....	46
Multigrid Algorithms on Parallel Processing Systems .....	46
Broyden Type Methods for the Solution of Linear and Nonlinear Systems of Equations .....	47
Multigrid as a Fast Poisson Solver .....	47
Approximation Problems which Arise in Connection with Polynomial Based Methods for Solving Linear Systems .....	47
Efficient Conjugate Gradient Type Methods for Solving Linear Systems with Complex Coefficient Matrices.....	48
Efficient Polynomial Preconditioners for Hermitian and General Non-Hermitian Matrices .....	48
Exact Solution of Linear Equations by Modular Arithmetic on a MIMD vs a SIMD Machine .....	49
<i>Parallel Algorithms</i>	
Time Dependent Partial Differential Equations .....	49
Higher Order Solution of the Euler Equations on Unstructured Grids Using Quadratic Reconstruction .....	50
Independent Time Stepping Methods.....	50
Solving the Shallow Water Equations on the Cray and Connection Machine.....	50
A Parallel Algorithm for Computing the Eigenvalues of Symmetric Tridiagonal Matrix.....	51

---

Ordered Fast Fourier Transforms on a Massively Parallel Hypercube Multiprocessor .....	51
Efficient Detection of a Weak Continuous Wave Signal with Linear Frequency Drift.....	51
Fast Fractional Fourier Transforms and Applications .....	52
Particle Simulation Methods on Parallel Computers.....	52
Optimal Moving Grids for Time Dependent Partial Differential Equations.....	52
Earth Mantle Convection on a Highly Parallel Hypercube .....	53
Iterative Methods for Markov Chain Modeling .....	53
Computing Radar Cross Sections .....	53
Parallel Sorting Algorithms .....	53
Efficient Solution Techniques for Problems Using Finite Element and Finite Volume Discretization Algorithms on Hypercube Architectures .....	54
<i>Architectures</i>	
Sparse Matrix Benchmark .....	54
Teraflops Computer .....	55
Assessment on Supercomputers .....	55
<i>Usability and Environments</i>	
The Paradigm Compiler .....	55
Solving Unstructured Problems on Massively Parallel Computers .....	56
A Massively Parallel Euler Solver for Unstructured Grids.....	56
Machine Independent Programming Languages for Highly Parallel Computers.....	57
Sparse Distributed Memory Simulator .....	57
Super-Workstations and Theoretical Chemistry .....	58
<i>Technology Transfer</i> .....	58
<b>Biographies</b> .....	61
<b>Appendices</b>	
Appendix 1 Publications by RIACS Staff.....	79
Appendix 2 Seminars, colloquia, and lectures .....	89
Appendix 3 Talks and presentations .....	93
Appendix 4 Other Activities of RIACS Staff .....	98





---

# Introduction

The Research Institute for Advanced Computer Science (RIACS) was established at the NASA Ames Research Center in June of 1983. RIACS is privately operated by the Universities Space Research Association (USRA), a consortium of 62 universities with graduate programs in the aerospace sciences, under a Cooperative Agreement with NASA. RIACS serves as the representative of the USRA universities at the Ames Research Center.

RIACS works as a partner with NASA to bring computer science into the US space program. This is accomplished by a research program emphasizing a small number of themes of strong relevance to NASA missions. In each thematic area, RIACS maintains a staff of full-time employees and a contingent of students and visitors from other research institutions, mostly universities. The entire staff work closely with NASA researchers.

The RIACS research themes support these grand challenges in science and engineering facing NASA:

Flying an airplane inside a computer.

Determining the chemical properties of materials under hostile conditions in the atmospheres of earth and the planets.

Sending autonomous machines on unmanned space missions.

Creating a one-world network that makes all scientific resources, including those in space, accessible to all the world's scientists.

Providing intelligent computational support to all stages of the process of scientific investigation from problem formulation to result dissemination.

Developing accurate global models for climatic behavior throughout the world.

Understanding very large data sets through visualization and automatic discovery.

In working with these challenges, we seek novel architectures and approaches to their use that exploit parallel and distributed computation and enable new functions that are beyond the current reach of computing machines. The investigation includes pattern computers as well as the more familiar numeric and symbolic computers, and it includes networked systems of resources distributed around the world and in space. We believe that the next significant results of computing research will come from interactions with other disciplines, and so all our projects are designed as partnerships with NASA.

This document reports our activities and accomplishments for the period from January 1, 1989, through June 30, 1990.

We at RIACS gratefully acknowledge the high level of technical and moral support we have received from NASA personnel at Ames and Headquarters.



---

# Overview

We at RIACS seek to contribute computing to solving the challenging problems being pursued by NASA at the frontiers of science. We do this by attracting the nation's most creative and productive computer scientists to the Ames Research Center to work on these problems and by working in close collaboration with NASA researchers in other disciplines.

The selection of problems for investigation is the foundation of our research. Our selection of problems is guided by our vision of scientific computing by the year 2000: computation, an integral part of scientific investigation, will require processing power far beyond what is possible with a single processor, will require functional power far beyond what is possible with a single node of a network, and will require cooperative power far beyond what is possible within a small group of investigators at a single site. In these three domains we see the required power arising from the collective, coordinated action of agents – tens of thousands of processors yielding computational power on the order of a teraflop and beyond, thousands of nodes connected by high-bandwidth networks yielding functional power on the order of thousands of functional compositions, and hundreds of people working together on common missions united by a high-speed network. We call these three the domains of data parallelism, distributed parallelism, and people parallelism. They are three interpretations of massive parallelism. Scientists must master these three domains to realize their vision of scientific computing by the end of the next decade.

RIACS people cannot address every significant problem suggested by the three domains of parallelism. We have, however, selected representative problems from these domains that are essential to NASA in their missions in science and technology. Our choices are reflected in our organization into three research areas:

*Learning Systems*

*Networked Systems*

*Parallel Systems*

In the area of *Learning Systems*, we seek to understand the design principles of machines that learn from training and experience, to provide solid mathematical grounding for the principles, and to build prototypes that demonstrate the principles. We use an operational definition of learning: an observer of the machine would say that the machine has, over a period of time, developed a new capacity for action that it did not have. Although this definition avoids the question of whether a machine that demonstrates learning is intelligent, we have taken inspiration from the structure of biological systems to design structures that can exhibit learning behavior. Within this group

we are pursuing two separate but related lines of investigation. The Sparse Distributed Memory (SDM) project is studying an associative memory structure that stores very long patterns (up to thousands of bits long) that represent encoded sensory data and associated actions, and can retrieve patterns when presented with partial cues. This memory is promising for automatic recognition of continuous speech, recognition of shapes in a visual field, and smooth motor control of a robot. The Bayesian Learning project is studying the applications of Bayesian statistics to inference of hypotheses about data. A program called AUTOCLASS II helped astronomers make new discoveries from infrared spectra data by grouping records into their most probable classes, and other applications of the theory will yield new approaches to understanding very large data sets that arise frequently through science. This work is being carried out with major sponsorship from the Computer Systems and Research, the Aerospace Human Factors Division, the Information Sciences Division of the Ames Research Center, and DARPA.

In the area of *Networked Systems*, we have been studying the concepts that will enable scientists to interact with each other and with resources in networks to carry out scientific computations and investigations, irrespective of the distances separating them. The resources include databases, servers, and instruments. One project within this area concerns the scientific computing environment, which has been examining tools that allow easy visualization of large data sets and easy construction of computations from resources attached widely around a network. A second project concerns networking within very high bandwidth nets, such as fiber optics, used in real-time control of components distributed around a network; these regimes will be at the core of networking in the Space Station. A third project concerns telescience applications, the conduct of scientific experiments that will be performed on the Space Station with assistance from ground-based investigators, astronauts, and mission control. A fourth project concerns global environment modeling that will bring together atmospheric, oceanographic, and other large-scale models into models that will permit forecasting of global environment conditions. This research is being carried out with major support from the Space Station and Advanced Technologies Office, the Life Sciences Division, the Information and Communications Systems Division, the Information Sciences Division, and the Space Life Sciences Payload Office of Ames Research Center.

In the area of *Parallel Systems*, we seek to understand the matches between novel, massively parallel computing architectures and scientific problems of interest to NASA, primarily computational fluid dynamics (CFD) and computational chemistry (CC). These investigations include the design and testing of algorithms capable of delivering the full power of thousands of processors to the problem, over a wide range of problem sizes -- e.g., algorithms that scale up almost linearly in the number of processors. These studies will ascertain whether a particular architecture has potential as a supercomputer and, if so, will develop efficient new algorithms that solve kernel problems in CFD and CC. Within the area, the Advanced Algorithms and Architectures project has concentrated on the basic linear algebra algorithms' underlying solution of partial differential equations; being in the inner loops of typical computations, these algorithms are the most likely source of the speedups achievable from parallel supercomputers. Basic libraries of these algorithms will be collected into a package called CFDTOOLS. The Center for Advanced Algorithms is a joint project involving NASA,

DARPA, and RIACS, studying the power of the Connection Machine 2 (CM-2) for scientific problems. This project is supporting deep studies of algorithms in image recognition, spectral analysis, molecular biology, text searching, and particle simulation of high altitude gases. These projects are being carried out with major support from the NAS Systems Division of the Ames Research Center and the Defense Advanced Research Projects Agency (DARPA).

With the exception of Director's "The Science of Computing" project, all RIACS projects have been assigned to divisions. About 80% of the funds for the projects come from tasks supported by NASA Ames research divisions, and about 20% from the core funds allocated by the Office of Aeronautics and Exploration Technology at NASA Headquarters. The core funds were used to support the time of the Director and portions of the time of others who engaged in seed projects, new starts, and strategic planning. The remainder of this section reports on the disposition of the core funds. The projects of the groups will be described in the three major sections following. Throughout this report, RIACS should be interpreted to mean "RIACS people acting in partnership with NASA people."

### *The RIACS Core*

The RIACS Core supported three types of activities:

Seed projects and short-term engagements to pursue discoveries made serendipitously in other projects.

Overall technical direction, leadership, and coordination of all the research activities of RIACS. This includes liaison with outside organizations relevant to NASA, especially the universities with whom collaborations are being conducted.

Computer workstation infrastructure for all staff.

The following are descriptions of the activities of the persons supported by the RIACS core for these activities.

#### **Peter Denning, Director.**

Denning was responsible for the overall technical leadership and management of RIACS. He maintained cognizance over RIACS projects and was responsible to keep them on track toward their objectives. He coordinated the activities of the scientific and administrative staff in support of the missions of RIACS. He assisted in raising resources for tasks. He worked with NASA people to define RIACS contributions to NASA initiatives such as the High Performance Computing Initiative (HPCI) and Global Change Technology Initiative (GCTI). He worked with NASA people to define specific ways in which RIACS personnel, computing resources, or expertise can be applied to NASA missions and problems. He engaged RIACS in seed projects of interest to NASA including the basic mathematical support for parallel sparse matrix computations, basic mathematical support for turbulence modeling, and basic technological support for collaboration over networks.

**Eugene Levin, Assistant Director.**

Levin's research time was devoted to identifying and specifying the requirements for graphics super-workstations designed for the scientific computing environment of NASA, in such domains as computational fluid dynamics and computation chemistry. He also performed research in computational chemistry, as reported later in this report.

**Barry Leiner, Senior Staff Scientist.**

Leiner served as Assistant Director for networked systems through March 1990, where he coordinated the technical projects within the Networked Systems group. He worked closely with NASA people in defining projects for the Telescience Testbed, a cooperative venture between ARC and the university community, and also to define functional requirements of advanced graphics in scientific workstation environments and other tools for the user interface such as X-windows. He initiated a new area of investigation in collaboration technologies, which are workstation-based tools that permit scientists in different locations to work on projects together, effectively without regard for the physical distance separating them. He served as a liaison between NASA activities in networking and other government activities, notably the Federal Research Internet Coordinating Committee (FRICC), the Internet Activities Board (IAB), and NSF networking advisory panels. Leiner's technical accomplishments are reported in the section for the Networked Systems.

**Michael Raugh, Chief Scientist.**

Raugh coordinated the technical projects within the Learning Systems area of RIACS. He worked with NASA people to define projects that will demonstrate the applicability of the Sparse Distributive Memory architecture to their domains, notably speech, visual shape recognition, robotics, and statistical prediction. He investigated the applications of SDM to cerebellar modeling and coordinated the SDM research with ARC projects in neuroscience. He was the principal liaison with outside industrial affiliates of the SDM project at Apple Computer, Hewlett-Packard Labs, and KLA Instruments. Raugh's technical accomplishments are reported in the section for the Learning Systems.

**Richard Sincovec, Assistant Director.**

Sincovec coordinated the technical projects within the Parallel Systems area of RIACS. He worked closely with NASA people to define projects that will demonstrate the applicability of highly parallel supercomputers containing thousands of processors to NASA problems and missions. He worked to establish a visiting scientist program to bring leading researchers in fluid dynamics to ARC for collaborations with NASA CFD researchers. He studied the programming language ADA as a practical means of programming highly parallel scientific computations. Sincovec's technical accomplishments are reported in the section for the Parallel Systems.

**Marjory Johnson, Senior Scientist.**

Core supported a portion of Johnson's time for strategic planning in the Networked Systems group.

**Robert Brown, Scientist.**

Core supported a portion of Brown's time for strategic planning in the Networked Systems group.

**Henry Sowizral, Scientist.**

Core supported a portion of Sowizral's time for a study of user services at national supercomputing centers.

**Tony Chan, Visiting Scientist (UCLA Mathematics Department).**

Core supported Chan's visits to Ames to work with the Parallel Systems group and NAS Systems Division on domain decomposition methods for parallel computers. Chan's technical accomplishments are reported in the section for the Parallel Systems.

**Roland Freund, Scientist.**

Core provided partial support for his studies of polynomial preconditioners for conjugate gradient methods and iterative methods on parallel machines in the Advanced Algorithms and Architectures project. A full report on those technical activities is given as part of the Parallel Systems report. This was a seed project while Freund sought and obtained full funding for his work.

**Niel Madsen, Visiting Scientist (Lawrence Livermore National Laboratory).**

Core supported Madsen's weekly visits to Ames to work with the Parallel Systems group and NAS Systems Division on time-dependent partial-differential equation solvers for parallel computers, and on finite volume methods for electromagnetic phenomena. Madsen's technical accomplishments are reported in the section for Parallel Systems.

**Robert Schreiber, Senior Scientist.**

Core provided partial support for his studies of continuous methods for nonlinear problems and of sparse-matrix computations on parallel computers in the Advanced Algorithms and Architectures project. A full report on those technical activities is given as part of the Parallel Systems report. This was a seed project while Schreiber sought and obtained full funding for his work.

**Carl Williams, Research Associate.**

Core supported a portion of Williams's time as he supported the research of the Parallel Systems group.

**William Behrman**, student (UC Berkeley).

Core supported Behrman's summer internship while Behrman worked with the Parallel Systems group and NAS Systems Division on continuous methods for non-linear problems. Behrman's technical accomplishments are reported in the section for Parallel Systems.

### *Distinguished Lecturer Series*

Core is supporting a lecture series called "Parallel Computing in the 90's" jointly with the NAS Systems Division. The speakers for the first part of 1990 were: Steve Lundstrom (January), Burton Smith (February), Ken Kennedy (March), Jack Dongarra (May).



---

# The Science of Computing

Peter J. Denning

This project is an ongoing effort to identify connections between the fundamentals of computer science and other scientific disciplines, mathematics, and engineering. The results have been published as a series of RIACS technical reports and articles in *American Scientist* magazine. The topics covered in 1985 were networks, parallel computing and its evolution, and the arbitration problem that arises when a circuit must distinguish between near simultaneous events. The topics covered in 1986 were expert systems, intelligent machines, random access memory, virtual memory, and electronic publishing. The topics covered in 1987 were network security, evaluation of supercomputer performance, multigrids and hypercubes, modeling of AIDS epidemiology, and cryptographic protocols for authentication in networks. The topics covered in 1988 were: the deadlock problem that arises when concurrent processes can stop and wait for signals, the unwanted limitations on design imposed by the traditional view that intelligence is problem-solving, computer viruses, speeding up parallel processing, memory protection, and computer models applied to the AIDS epidemic.

Ten topics were studied in 1989-1990:

- **Massive Parallelism in the Future of Science**

Massive parallelism appears in three domains of action of concern to scientists, where it produces collective action that is not possible from any individual agent's behavior. In the domain of data parallelism we will design computers comprising very large numbers of processing agents, one for each data item in the result; these agents collectively can solve problems thousands of times faster than current supercomputers. In the domain of distributed parallelism we will design computations comprising large numbers of resource attached to the world network; the network will support computations far beyond the power of any one machine. In the domain of people parallelism, we will design collaborations among large groups of scientists around the world who participate in projects that endure well past the sojourns of individuals within them; computing and telecommunications technology will support the large, long projects that will characterize "big science" by the turn of the century. Scientists must become masters in these three domains during the coming decade. [*American Scientist* 77, No. 1 (January-February 1989). Also RIACS TR88.33.]

- **The Internet Worm**

In November 1988 a worm program invaded several thousand UNIX-operated Sun workstations and VAX computers attached to the Research Internet, seriously disrupting service for several days but damaging no files. An analysis of the worm's decompiled code revealed a battery of attacks by a knowledgeable insider, and demonstrated a number of security weaknesses. The attack occurred in an open network, and little can be inferred about the vulnerabilities of closed networks used for critical operations. The attack showed that password protection procedures need review and strengthening. It showed that sets of mutually trusting computers need to be carefully controlled. Sharp public reaction crystalized into a demand for user awareness and accountability in a networked world. [*American Scientist* 77, No. 2 (March-April 1989). Also RIACS TR89.03.]

- **Bayesian Learning**

In 1983 and 1984, the Infrared Astronomical Satellite (IRAS) detected 5,425 stellar objects and measured their infrared spectra. In 1987 a program called AUTOCLASS used Bayesian inference methods to discover the classes present in these data and determine the most probable class of each object, revealing unknown phenomena in astronomy. AUTOCLASS has rekindled the old debate on the suitability of Bayesian methods, which are computationally intensive, interpret probabilities as plausibility measures rather than frequencies, and appear to depend on a subjective assessment of the probability of a hypothesis before the data were collected. Modern statistical methods have, however, recently been shown to also depend on subjective elements. These debates bring into question the whole tradition of scientific objectivity and offer scientists a new way to take responsibility for their findings and conclusions. [*American Scientist* 77, No. 3 (May-June 1989). Also RIACS TR89.12.]

- **Sparse Distributed Memory**

Sparse distributed memory was proposed by Pentti Kanerva as a realizable architecture that could store large patterns and retrieve them based on partial matches with patterns representing current sensory inputs. This memory exhibits behaviors, both in theory and in experiment, that resemble those previously unapproachable by machines -- e.g., rapid recognition of faces or odors, discovery of new connections between seemingly unrelated ideas, continuation of a sequence of events when given a cue from the middle, knowing that one doesn't know, or getting stuck with an answer on the tip of one's tongue. These behaviors are now within reach of machines that can be incorporated into the computing systems of robots capable of seeing, talking, and manipulating. Kanerva's theory is a break with the Western, rationalistic tradition, allowing a new interpretation of learning and cognition that respects biology and the mysteries of individual human beings. [*American Scientist* 77, No. 4, (July-August 1989). Also RIACS TR89.22.]

- **Worldnet**

The expanding use of powerful workstations coupled to ubiquitous networks is transforming scientific and engineering research and the ways organizations around the world do business. By the year 2000, few enterprises will be able to succeed without

mastery of this technology, which will be embodied in an information infrastructure based on a worldwide network. A recurring theme in all the discussions of what might be possible within the emerging Worldnet is people and machines working together in new ways across distance and time. This essay reviews the basic concepts on which the architecture of Worldnet must be built: coordination of action, authentication, privacy, and naming. Worldnet must provide additional functions to support the ongoing processes of suppliers and consumers: help services, aids for designing and producing subsystems, spinning off new machines, and resistance to attack. This discussion begins to reveal the constituent elements of a theory for Worldnet, a theory focused on what-people-do-with-computers rather than on what-computers-do. [*American Scientist* 77, No. 5 (September-October 1989). Also RIACS TR89.26.]

- **The Arpanet after Twenty Years**

The ARPANET began operation in 1969 with four nodes as an experiment in resource sharing among computers. It evolved into a worldwide research network of over 60,000 nodes, influencing the design of other networks in business, education, and government. It demonstrated the speed and reliability of packet-switching networks. Its protocols have served as the models for international standards. And yet the significance of the ARPANET lies not in its technology, but in the profound alterations networking has produced in human practices. Network designers must now turn their attention to the discourses of scientific technology, business, education, and government that are being mixed together in the milieu of networking, and in particular the conflicts and misunderstandings that arise from the different worldviews of these discourses. [*American Scientist* 77, No. 6 (November-December 1989). Also RIACS TR89.38.]

- **Stopping Computer Crimes**

Two new books about intrusions and viruses remind us that attacks against our computers on networks are the actions of human beings. Cliff Stoll's book about the wily hacker who spent a year beginning August 1986 attempting to use the Lawrence Berkeley Computer as a stepping-stone for access to military secrets is a spy thriller that illustrates the weaknesses of our password systems and the difficulties in compiling evidence against a hacker engaged in espionage. Pamela Kane's book about viruses that attack IBM PC computers shows that viruses are the modern version of the old problem of Trojan horse attack. It discusses the most famous viruses and their countermeasures, and it comes with a floppy disk of utility programs that will disinfect your PC and thwart future attack. [*American Scientist* 78, No. 1 (January-February 1990). Also RIACS TR88.47.]

- **Is Thinking Computable?**

Strong AI claims that conscious thought can arise in computers containing the right algorithms even though none of the programs or components of those computers "understands" what is going on. As proof, it asserts that brains are finite webs of neurons, each with a definite function governed by the laws of physics; this web has a set of equations that can be solved (or simulated) by a sufficiently powerful computer. Strong AI claims the Turing test as a criterion of success. A recent debate in *Scientific American* concludes that the Turing test is not sufficient, but leaves intact the under-

lying premise that thought is a computable process. The recent book by Roger Penrose, however, offers a sharp challenge, arguing that the laws of quantum physics may govern mental processes and that these laws may not be computable. In every area of mathematics and physics, Penrose finds evidence of non-algorithmic human activity and concludes that mental processes are inherently more powerful than computational processes. [*American Scientist* 78, No. 2 (March-April 1990). Also RIACS TR90.02]

- **Changing the Guard**

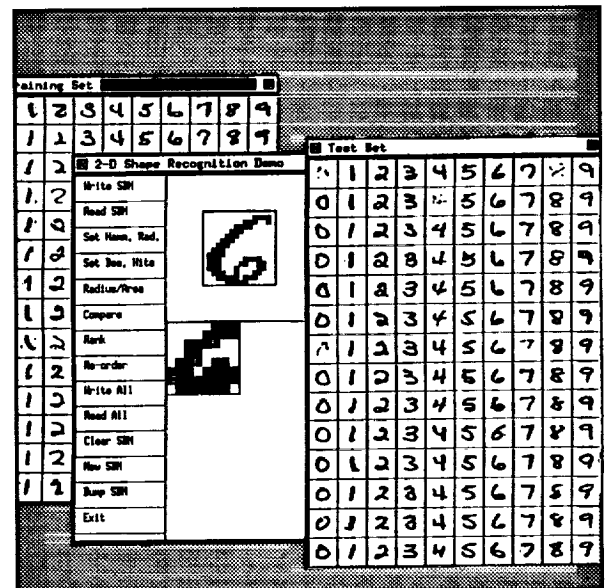
Columnist Denning bids farewell to the editorial staff of the *American Scientist* magazine on the occasion of their departure. He discusses his own interests and experience with the magazine. He looks behind the scenes at how the staff is organized, how they interact with the board of editors, and what fundamental values they embody. He salutes their leadership among editors and publishers in science. [*American Scientist*, No. 3 (May-June 1990). Also RIACS TR90.17.]

- **Information Technologies for Astrophysics Circa 2000**

Denning prepared this as a position paper for presentation before a NASA workshop in May 1990 that was planning a research programs for information systems in astrophysics for the rest of the 1990s. He reported on trends on miniaturization, multiprocessing, software technology, networking, databases, graphics, pattern computation, and interdisciplinary studies. He challenged three paradigms that shape our thinking about information systems. "Saving all the bits" holds that all bits from instruments and massive computations must be saved because the cost of recovering them is too high, or because some discovery might be lost to science. "Technology off the shelf" holds that NASA can obtain all networking and data storage technologies it requires from commercial vendors. "Linear model of innovation" holds that innovation begins with a discovery and passes successively through the pipeline of research, development, production, and market on the way to the customer. He argues instead that there are important cases where machines can make discoveries that would otherwise be missed by humans and that such machines can reduce the information load to be within the capacity of our networks and storage systems; that NASA has unique requirements that demand a research program in information systems; and that most innovation occurs in ongoing cycles of improvements. [To be issued as a RIACS technical report in August 1990.]

Candidates for study in the future include synchronization of clocks in distributed systems, saving all the bits of massive data sets, discrete event simulation on parallel machines, and beating the Cray YMP computer.

# Learning Systems



From the display panel of a two-dimensional shape recognition experiment. The SDM-based system 'learns' the handwritten numerals of a one-hundred-forty word training set in less than two minutes on standard workstation.



---

# Learning Systems

One of NASA's grand challenges is to build autonomous machines and systems that are capable of learning to perform tasks too tedious, or in places too remote and too hostile, for humans. The goal of the Learning Systems of RIACS is to find new approaches to autonomous systems based upon sound mathematical and engineering principles, to understand how information processing is organized in animals, and to test the applicability of these new approaches to the grand challenge. The research program includes the development of theory, implementations in software and hardware, and explorations of potential areas for applications. The Learning Systems consists of two projects that pursue these goals, the Sparse Distributed Memory Project and the Bayesian Learning Project.

## The Sparse Distributed Memory Project

The Sparse Distributed Memory (SDM) Project began at RIACS in late 1985 as an effort to study the applicability of sparse distributed memory to the detection, identification, and manipulation of physical objects. In our original five-year plan, we proposed to develop theory, to develop simulators, and to choose problems that could lead to the understanding of how to apply SDM in chosen broad application areas. The work reported here is a culmination of that five-year program. The work is also a springboard for initiatives to apply SDM to specific problems of importance to NASA. Nine related SDM projects are reviewed below:

- SDM Alternative Designs
- Speech Recognition
- Digital Prototype of SDM
- Cerebellar-Model Associative Memory
- Fine Discrimination based on Distinctive Features
- Encoding Features for SDM Image Recognition
- Two-Dimensional Shape Recognition
- Large-Scale Simulation on the CM-2 Connection Machine
- Weather Prediction.

SDM is a new kind of associative random-access memory that uses very large binary patterns -- hundreds to thousands of bits in length -- as both addresses and data. The

patterns are well suited to represent encoded information derived from sensors, for example, from the visual, auditory, or tactile sensors of a robot. SDM affords an architecture that can be used to study systems that learn from experience, and, in particular, it provides a unified approach to problems in shape recognition, speech recognition, systems control, and automatic processing of very large data sets. Special to the SDM Project is the extent to which our engineering designs and mathematical models are driven by the desire to understand how information is processed in animals. The original design of SDM was developed by Dr. Pentti Kanerva, the Project Principal Investigator. The basic theory is set forth in his book *Sparse Distributed Memory* [MIT Press, 1988].

In the ordinary random access memory (RAM), common in today's computers, addresses are typically 16-to-32 bits in length. In a RAM each item of data is stored in one memory location, each address points to an actual hardware location in the memory, and each stored item is recalled exactly. In contrast, patterns used as addresses in an SDM are much longer, the data are distributed over many locations, a sparse subset of addresses point to actual hardware locations, and information read from the memory is a statistical reconstruction that may not be identical to what was written into the memory. When writing a pattern at an address in SDM, the pattern is added to existing information at each of many nearby memory locations. When reading from an address in the memory, information stored at nearby memory locations is *pooled* and *thresholded* for output.

The sparse distributed memory can either be interpreted as a special kind of neural network or as a random-access memory. Its potential utility is a result of its statistical properties and of other factors: (1) A large pattern representing an object or a scene or a moment of experience can encode a large amount of information about what it represents. (2) This information can serve as an address to the memory, and it can also serve as data. (3) The memory can interpolate and extrapolate from existing data and is fault tolerant. (4) The memory is also noise tolerant -- the information need not be exact. (5) The memory can be made very large, and large amounts of information can be stored in it. (6) The memory can store long sequences of patterns and can *predict* (recall) the remaining portion of a sequence when prompted by an earlier segment of the sequence. (7) The architecture is inherently parallel, allowing large memories to be fast. (8) The mathematical theory is clear-cut and is well understood. (9) Learning is fast -- only a small number of training cycles are necessary. [For an additional overview, see P. J. Denning, "Sparse Distributed Memory," *Am. Sci.* 77 No. 4, July-August '89, pp. 333-335.]

The SDM Project staff combines graduate-level expertise in mathematics, statistics, artificial intelligence, electrical engineering, chemistry, physics, and philosophy. The Project scientists during 1989 were Douglas Danforth, Louis Jaeckel, Pentti Kanerva, and David Rogers. Michael Raugh was Project Manager. The Project enjoyed the summer visits of university colleagues, Professor Richard Schori of OSU and his student Isa Jubran, research associate Andre Lamothe of San Jose State University, and the remaining nine months of a one-and-a-half-year visit by graduate student, Bruno Olsahusen (Caltech) as research associate. Dr. Charles C. Jorgensen was a Visiting Scientist for the first six months of 1990, with special support from the Ames Information Sciences Division and from Dr. Dan Palumbo of the NASA Langley Research Center.



Professor Kimmo Kaski of the University of Tampere visited the Project early in 1990. The Project received additional cooperation and support from the Information Sciences Division (graduate student, Coe Miles) and the Human Factors Division (Andrew Watson of Ames Research Center). Pentti Kanerva was absent from the Project while on a six-month visit to the University of Tampere in Finland.

The SDM Project is funded by NASA, with additional support from Apple Computer Corporation (equipment gift), DARPA (support for work on the Connection Machine), and Hewlett-Packard Laboratories (participation of Egon Loebner\*, Counselor for Science and Technology at HP Labs of Palo Alto).

### A. Objectives of the Project

- I. *Architectural studies*: Develop and analyze SDM designs well-suited for particular types of real-world data.
- II. *Sensory encoding for SDM*: Develop techniques for encoding real-world data for use in SDM.
- III. *Applications and technology transfer*: Test the applicability of SDM for broad application areas and promote cooperative efforts with NASA scientists for experimental applications in areas of importance to NASA.

### B. Strategies for achieving the objectives

Our strategy to date, in fulfillment of our five-year charter, has been to develop the theory of SDM and, through analysis and simulations, test the applicability of SDM to selected problems in broad application areas, such as, shape recognition, speech recognition, systems control, and automatic processing of very large data sets. Every application falls into two interrelated parts: 1) appropriate design of the memory and its learning algorithm for the problem, and 2) encoding of the data for the problem; each a challenging problem in its own right. Thus, there is an inherent ambiguity in assigning a project to categories I, II, and III. Suffice it to say that any projects classified as *application* entails important aspects of *architecture* and *encoding*.

In the review period we advanced far in learning how to apply SDM to broad problem areas and have actively published our efforts. To make the work widely available, we have published fourteen technical reports, seven of which will have (or will) appear as conference proceedings. We also presented a session at the 34th IEEE Computer Society International Conference in San Francisco (COMPCON Spring 89) and led an all-day workshop at the IEEE Neural Information Processing Systems Conference at Keystone. Some interesting work on SDM done at other institutions includes a patent application for a VLSI implementation of SDM by Derek Smith and associates at Texas Instruments, and publications from the Department of Engineering at Cambridge University (Prager, et. al.) and from JPL/Caltech (Anderson):

Prager, R.W., Clarke, T.J.W., and Fallside, F. "The Modified Kanerva Model: Results for Real Time Word Recognition," *Proceedings IEEE First International Conference on Artificial Neural Networks*, pp. 105-109, 16th - 18th October, 1989. London, U.K.

Anderson, C., "A Conditional Probability Interpretation of Kanerva's Sparse Distributed Memory," *Proceedings of the IJCNN Meeting* (International Joint Conference on Neural Networks – sponsored by the IEEE and International Neural Network Society), Washington, D.C., June 18-22, 1989, Vol. I, pp. 415-418.

Prager, R. W. and Fallside, F. "The modified Kanerva model for automatic speech recognition," *Computer Speech and Language*, 3, (1989), 61-81.

In the future, we will continue with theoretical studies, for there is much yet to be learned. But we will also emphasize applications of direct interest to NASA. Having succeeded in our initial validations, we now look forward to more direct interactions with NASA scientists.

## C. Accomplishments

### *Architectural Studies*

- **SDM Alternative Designs**

Louis Jaeckel

Louis Jaeckel reported a broad class of alternative designs for SDM. Viewed as a kind of random access memory with a very large address space, SDM is characterized by the fact that each addressing operation activates not just one hardware storage location (*hard location*) but many such locations. This is true for both read and write operations. For example, in Kanerva's original design (Kanerva, MIT Press, 1988), the activated locations are determined by comparing the reference address to the address of each hard location; if the Hamming distance between the reference address and a hard location is smaller than a criterion value, then that hard location is activated, otherwise it is not. The expected number of activated hard locations for any given reference address is often a large number – 1,000 out of a million in Kanerva's original model. But Hamming distance is not the only way to distribute data in an SDM. Jaeckel has proposed a broad class of designs, different from the original one, that have the desirable properties of being relatively easy to analyze mathematically and easy to implement in hardware as well as software. For example, in the *selected-coordinates* design, each hard location is determined by, first, selecting ten coordinates of the address at random, then, second, assigning a random bit value for each of the ten coordinates. A given hard location is activated if the reference address matches the selected coordinates of the hard location in all ten of the assigned bits. Jaeckel has broadened the class of alternative designs to include a hyperplane design and has shown how these designs form a spectrum of related designs, with Kanerva's original design at one extreme and the selected-coordinate design at the other. A patent application was filed, covering the spectrum of alternative designs.

[See Jaeckel, Louis A., "A Class of Designs for a Sparse Distributed Memory," RIACS TR89.30, July 1989, and Jaeckel, Louis A., "An Alternative Design for a Sparse Distributed Memory," RIACS TR89.28.]

- **Speech Recognition**

Douglas Danforth

The original theory of SDM was based upon the hypothesis that data are distributed randomly throughout the address space. The assumption of randomness permitted Kanerva to gain insight from rigorous mathematical deductions. To extend the theory to real-world applications, it is necessary to deal with data that are not randomly distributed. Douglas Danforth has used discrete-speech recognition as a probe to reveal fundamental factors that affect the performance of SDM for highly correlated, real-world data.

In a sequence of progressive experiments involving training and testing on the spoken words for the ten digits from zero to nine, Danforth found that *orthogonal codes* (binary patterns that differ in half their bits) for word labels (i.e., desired output patterns) substantially improve performance -- from 49.6% for ASCII labels to 81.2% for orthogonal labels. Next, Danforth noted that the tenets of the basic SDM model were grossly violated by the speech data -- the mean of the encoded speech data was 11 standard deviations away from the mean of random data. To better deal with these correlated data, Danforth compared an SDM, whose location addresses were chosen at random, with an SDM whose location addresses were the encoded patterns from speech. A total of 850 patterns from discrete speech were used for the location addresses. Performance jumped to 92.1%. Next, to improve the performance of a read operation, Danforth modified the pool-and-threshold operation performed in the basic SDM. Instead of thresholding the counter sums, he used the sums directly to obtain inner products of the sums with the class labels and chose the class corresponding to the maximum inner product. Performance rose to 94.6%. This technique selects the class of maximum frequency. Another gain was obtained by slightly moving the address of each hard location to increase the signal-to-noise ratio at each location. In this way, Kanerva's original design was modified incrementally to improve performance to 97.5%. Finally, Danforth experimented with Jaeckel's selected-coordinate design, in which only three bits of the input are examined by each hard location and the decision to activate is determined by whether those bits match exactly the corresponding bits of the hard location. This gave a 97.1% recognition score. Incorporating a learning rule proposed by Richard Prager and associates at Cambridge University (which Danforth calls the *polarity rule*), the recognition score rose to 99.3%. The polarity rule adjusts the counters of activated memory locations just enough to produce the correct output if the same input were presented again. This adjustment sometimes increases and sometimes decreases a counter's value.

Theoretical work by Danforth supports the heuristic polarity rule in that, for both the basic SDM model and the selected-coordinate design, it is possible to derive an inverse. By using the inverse for writing and standard pooling of activated locations for reading, the signal-to-noise ratio of the memory can be increased. The form of the inverse specifies a write rule that is reminiscent of edge enhancement techniques found in animal vision (on-center off-surround).

[See Danforth, Douglas G., "An Empirical Investigation of Sparse Distributed Memory Using Discrete Speech Recognition," RIACS TR90.18.]

[See also Danforth, Douglas G., "Speech Transcription Using Sparse Distributed Memory," RIACS Memorandum 89.02, 1989.]

- **Digital Prototype of SDM**

Michael Flynn/Stanford, Brian Flachs/Stanford, Andrew Zimmerman/Stanford, Pentti Kanerva, Andre Lamothe

In order to demonstrate the engineering feasibility of SDM, NASA has supported the development of a digital prototype at Stanford University's Computer Systems Laboratory under the direction of Michael Flynn. The work has proceeded with the cooperation of the Sparse Distributed Memory Project at RIACS; Pentti Kanerva consulted on architecture, and Andre Lamothe provided software support.

A four-fold Sparse Distributed Memory system was completed and is now available for use. This is a modified version of a one-fold system developed in 1988. The new system features five 68030 processors, special purpose matching hardware, and a library of primitives for C. It handles up to four times as many counter operations as its predecessor. This added work load necessitated a parallel solution to maintain system performance. To this end, five FORCE CPU-30 VME cards were installed, which offer two features facilitating interprocessor communication: dual ported memory and a VME broadcast mode. Once the hardware was chosen, the software effort was divided between a special purpose system run by the SDM processors and tools available to users on the host system, currently a SUN 360. Performance considerations were important for both efforts.

The internal structure of the SDM system is guided by the natural divisions of the physical communication methods. Although a few SDM commands are executed on the host system, the host merely sends the command and its parameters over the SCSI bus where the majority of the functions are directly implemented by the Command Module. The Command Module processes any commands related to either the Address Module, the tag lists, or any of the status registers. If, however, the operation requires manipulation of the fold memories, a message is broadcast simultaneously starting all four fold processors. Each processor is responsible for processing of a data word. This partitioning allows SDM to process single fold operations much faster than a system that pairs a processor to a fold. The primitive library communicates with the SDM system via a SCSI bus. When the communication is initiated, the SDM returns a version identifier. If the library is incompatible with SDM, an error is returned.

We plan to continue cooperation with Stanford in evaluation and analysis of the digital prototype and to assist in efforts to derive new architectural approaches.

[See Flynn, Michael J., Pentti Kanerva, and Neil Bhadkamkar, "Sparse Distributed Memory: Principles and Operation," RIACS TR89.53. (Published concurrently at Stanford University's Computer Systems Laboratory as Technical Report CSL-TR89.40.)]

- **Cerebellar Modeling**

Egon Loebner/HP Labs, Coe Miles/NASA, David Rogers

This ongoing work involves Egon Loebner of Hewlett-Packard Laboratories,\* Coe Miles of NASA Ames Information Sciences Division, and David Rogers of RIACS. An extensive review of current neuroscience literature was completed to learn about

motor control and the processing of sensory data in mammals. We have been interested in the cerebellum for two reasons. First, the cerebellum in mammals coordinates all modes of sensory inputs (except olfaction) with motor effector outputs. It is also a highly regular structure, perhaps more so than any other tissue in the brain. In a sense then, it appears to be a general computer for coordinating hundreds of reflex arcs robustly, in real time. Any elucidation of the relationship between structure and function of the cerebellum will have important implications for applications of computer science to advanced control problems. Second, sparse distributed memory is similar to the cerebellum -- Kanerva has drawn attention to the similarity between SDM and earlier cerebellar models of associative memory by Marr and by Albus.

Coe Miles, a doctoral candidate at Santa Clara University, has begun a thesis project developing a software model of the cerebellum. This work will be reported in 1990.

[See articles by Kanerva and by Loebner in Proceedings of IEEE COMPCON 89, Session 21B, Cerebellar Models of Associative Memory, San Francisco, CA, February 27 - March 3, 1989, Washington, D.C., IEEE Computer Society Press, 1989, or in Raugh, Michael R., Editor, *Cerebellar Models of Associative Memory: Three Papers from IEEE COMPCON Spring '89*, also RIACS TR89.11.]

[See also Loebner, Egon, Coe Miles-Schlichting, and David Rogers, "Engineering Aspects of the Cerebellum: A Collaboration," RIACS Memorandum 89.04, 1989.]

## *Sensory Encoding*

- **Fine Discrimination Based on Distinctive Features**

Pentti Kanerva, Isa Jubran/OSU

Pentti Kanerva described a multipass method for discriminating between nearly identical patterns based on a small number of distinctive features, for example, the letters "P" and "R". The sparse distributed memory and many other distributed memory models lump similar patterns together and, in normal operation, cannot in one read operation distinguish between patterns that differ in only a small number of bits. Multipass reading finds first a subset of training patterns similar to the present test pattern, and identifies a subset of bits (the distinctive features) at which the similar patterns differ. The second pass uses those bits to identify the test pattern or to find a subset of patterns and the corresponding distinctive features for the third pass. Each pass of this recursive process reduces the subset of possible patterns and the number of distinctive features, allowing discrimination based on even one distinctive feature. Isa Jubran, a graduate student from Oregon State University, began a study of the method as a summer visitor at RIACS and is continuing the work at the University, with results to be reported in 1990.

- **Encoding Features for SDM Image Recognition**

Louis Jaeckel

Louis Jaeckel reported ongoing work on encoding simple visual images for use with a sparse distributed memory. He considers the class of images that are composed of arcs and line segments. The class includes line drawings of characters, such as letters

of the alphabet. He describes a method of representing a segment or an arc by five numbers in a *continuous* way; that is, similar arcs have similar representations. He also gives methods of encoding these numbers as bit strings in an approximately continuous way. In recent work he describes two possible implementations of an SDM for which encoded images will serve as both addresses to the memory and as data to be stored in the memory. In the first model, an image is encoded as a 9072-bit string, which may be used as a read or a write address. Another encoding scheme, in which an image is encoded as a 256-bit string, may be used with either model as data to be stored in the memory, but not as an address. Since an image can be approximately recovered from this encoding, it is possible to do a sequence of read operations, using the result of each read as the next read address. In the second model, an image is not converted to a vector of fixed length. Instead, Jaeckel gives a rule for determining which memory locations are to be activated in response to an encoded image. The rule treats pieces of an image as an unordered set. Jaeckel performed experiments with a small-scale simulation of the second model.

[See Jaeckel, Louis A., "Recognition of Simple Visual Images Using a Sparse Distributed Memory: Some Implementations and Experiments," RIACS TR90.11.]

[See also Jaeckel, Louis A., "Some Methods of Encoding Simple Visual Images for Use with a Sparse Distributed Memory, with Applications to Character Recognition," RIACS TR89.29.]

### *Applications and Technology Transfer*

- **Two-Dimensional Shape Recognition**

Bruno Olshausen, Andrew Watson/NASA, Pentti Kanerva

Bruno Olshausen, working with NASA scientist Andrew Watson and Pentti Kanerva, developed and demonstrated an SDM-based system for recognizing two-dimensional shapes using an image preprocessor based upon Watson's cortex *transform*. In one experiment (to be reported early in 1990), a set of images were injected with gaussian noise, preprocessed with the cortex transform, and then encoded into bit patterns. The various spatial-frequency bands of the cortex transform were encoded separately so that they could be evaluated based on their ability to properly cluster patterns belonging to the same class. The results of the study indicate that by simply encoding the low-pass band of the cortex transform, a very suitable input representation for the SDM can be achieved. In another experiment using the same encoding method and low-pass filter, Olshausen built and demonstrated a system for recognizing characters from the US Postal Service Database of Handwritten Zipcodes. After training the system on 140 random characters, the system was tested on 140 different random characters. It correctly labelled 86% of the test characters. The importance of this demonstration is that the method is very general in the sense that no special features of written characters were used to improve performance. Kanerva has provided a related theoretical discussion of SDM-based shape recognition, using a very general contour-map representation of two-dimensional shapes (Kanerva, 1990).

[See Olshausen, B.A., and Watson, A.B., "The Cortex Transform as an Image Processor for Sparse Distributed Memory: An Initial Study," RIACS TR90.04.]

[See also Kanerva, Pentti, Bruno Olshausen, "Two-Dimensional Shape Recognition Using Sparse Distributed Memory," RIACS Memorandum 89.03, 1989.]

[See also Kanerva, Pentti, "Contour-Map Encoding of Shape for Early Vision," RIACS TR 90.05, February 1990 (see also under Conference Proceedings).]

- **Large-scale Simulation on the CM-2 Connection Machine**

David Rogers

David Rogers continued the development of an SDM simulator on the CM-2 Connection Machine. An early version of the CM-2 software was translated from LISP to C/Paris and integrated into a unified front-end, to permit use of the same command sequences to operate an SDM on the Connection Machine, on the Stanford hardware, and on the local host. This software is ready for alpha-test release. Timings were taken of this new simulator. Extrapolations for a full 64K CM-2 imply that a 64K location memory would perform approximately 166 addressing operations per second, 30 write operations/second, and 4 read operations per second; a 1M location memory would perform approximately 17 addressing operations per second, 7 write operations/second, and 1.3 read operations/second. A technical report, documenting installation procedures and operation of the simulator, has been drafted and will be published in 1990.

[See Rogers, David, "Kanerva's Sparse Distributed Memory: An Associative Memory Algorithm Well-Suited to the Connection Machine," in *Proceedings of a Conference on Scientific Applications of the Connection Machine*, NASA Ames Research Center, CA, September 12-14, 1988, book published by World Scientific Publishing, 1989.]

[See also Rogers, David, "Kanerva's Sparse Distributed Memory: An Associative Memory Algorithm Well-Suited To The Connection Machine," *International Journal of High Speed Computing*, vol. 1, No. 2, pp. 349-365, 1989.]

[See also Rogers, David, "Kanerva's Sparse Distributed Memory: An Associative Memory Algorithm Well-Suited To The Connection Machine," RIACS TR 88.32.]

- **Weather Prediction**

David Rogers

David Rogers has combined Kanerva's SDM with Holland's genetic algorithms in what he calls *genetic memory*. The genetic memory is an outgrowth of Rogers' earlier work on the use of SDM for statistical prediction.

SDM is an associative-memory model based on the mathematical properties of high-dimensional binary address spaces. The genetic algorithms are a search technique for high-dimensional spaces inspired by evolutionary processes of DNA. The genetic memory uses a genetic algorithm to dynamically reconfigure the addresses of physical storage locations in an SDM to reflect correlations between the stored addresses and data.

Initial testing of the genetic memory has been done using approximately sixty-thousand samples of weather data collected over a twenty-five year period at a meteorological station in Perth, Australia. A 256-bit SDM was simulated, using 1,000 randomly selected addresses for the storage locations, and each storage location was

assigned just one counter. For each sample of weather data, fifteen meteorological features, such as wind direction and speed, wet-bulb temperature, barometric pressure, and dew point, were coded and incorporated into a bit pattern that served as an address. At that address in the memory one bit was written indicating whether it rained in the four-hour period immediately following the time of the sample. After each dozen of such inputs, all the current memory locations were read to determine which locations recorded a higher-than-average density of rainfall. Then two locations with relatively high density were selected and *bred*, i.e., a new 256-bit address was created by mixing random complementary segments from the addresses of the *breeding* pair. The resulting address was assigned to a location that registered a near-average density of rainfall. In other words, locations that did not score well as rain predictors were replaced by the offspring of those that did score well. Thus, after one pass over all sixty-thousand data samples, the memory reconfigured itself five thousand times. The result was that several locations were generated automatically that correlated well with rainfall. Because the addresses of locations are stored in the SDM, it is straightforward to examine the SDM array of addresses, thence to see explicitly which bits (and meteorological features) corresponded to rain. As an example, the genetic memory found that rain was likely to occur when the following combination of features exist: high cloud cover, southerly wind, January-February, pressure dropping. Rogers estimates that the weather data could have been processed on the Connection Machine in less than fifteen minutes. This work will be reported in 1990.

It appears that genetic memory is well suited for finding correlations in large data sets involving large numbers of independent variables. We plan to test genetic memory on NASA data sets, such as satellite spectral data, to help scientists detect and study correlations with phenomena of interest, such as terrestrial methane release, to help in predicting greenhouse effects.

[See Rogers, David, "Statistical Prediction with Kanerva's Sparse Distributed Memory," RIACS TR89.02.]

[See Rogers, David, "Predicting Weather using a Genetic Memory: A Combination of Kanerva's Sparse Distributed Memory and Holland's Genetic Algorithms," in D.S. Touretsky (ed.), *Advances in Neural Information Processing Systems 2*, San Mateo, California: Kaufmann, 1990.]

[See also Rogers, David, "Weather Prediction using a Genetic Memory," RIACS TR90.06.]

- **Applications of High-Dimensional Associative Memories to Robot Control**  
Charles C. Jorgensen

Jorgensen explored the suitability of two varieties of distributed-memory neural networks as trainable controllers for a simulated robotics task. The task required that two cameras observe an arbitrary target point in space. Coordinates of the target on the camera image planes are passed to a neural controller that must learn to solve the inverse kinematics of a manipulator with one revolute and two prismatic joints. Two types of network design were evaluated.



The first type of network, called a *radial-basis sparse distributed memory*, approximates functional mappings as sums of multivariate gaussians centered around previously learned patterns. New patterns are categorized using gaussian distances from learned examples with training patterns stored as high-dimensional input addresses in a sparse distributed memory architecture.

The second type of network involved variations of adaptive vector quantizers or self-organizing maps. Random n-dimensional points are given local connectivities. They are then exposed to training patterns and readjust their locations based on a nearest neighbor rule. The winning point and its neighbors are dragged differentially toward the new pattern, adjusting to minimize the elastic global energy of the network. The result is a network that adaptively forms an interpolating n-dimensional surface over the density of the training sample set. A new learning rule is proposed, called the *proportional winner* rule, which dramatically simplifies problems in learning rate and radius scheduling. A new network called an *infolding net* is presented, which has advantages of a self-organizing map with superior learning performance and potential for real-time control.

Both approaches were evaluated based on their ability to interpolate manipulator joint coordinates for simulated arm movement while simultaneously performing stereo fusion of the camera data. Comparisons were made with classical k-nearest neighbor pattern recognition techniques. A procedure for application testing with real hardware was described. This work has been reported.

[See Jorgensen, Charles C., "Development of a Sensor Coordinated Kinematic Model for Neural Network Controller Training," RIACS TR90.28.

See also Jorgensen, Charles C., "Distributed Memory Approaches for Robotic Neural Controllers," RIACS TR90.29.]

\*On December 30, 1989, Dr. Egon Loebner died of cancer. Dr. Loebner was one of the original pioneers in the field of neurocomputing. He developed an electro-optic implementation of the frog's retina in the early 1960s. His knowledge of biological and technological evolution as well as his grasp in such diverse fields as engineering, biology, physics, chemistry, and psychology made him the inspirational leader of the Cerebellar Research Group. We especially wish to thank Hewlett-Packard Laboratories of Palo Alto for support of Dr. Loebner's participation in this project.

## The Bayesian Learning Project--(AutoClass)

Bayesian learning is a form of statistical learning that is particularly good at finding patterns in noisy data. In addition, the Bayesian theory can guarantee that the patterns found represent a real effect operating in the data, and not an artifact of the data analysis. In other words, the Bayesian method is particularly good at distinguishing signal from noise in very noisy data. It is also well suited for using prior knowledge and statistics to guide the selection of patterns among competing alternatives. While the general Bayesian theory has been around for many years, the discovery of its implications in particular domains is still largely unexplored.

Traditional classification methods, such as cluster analysis, have always had a problem distinguishing signal from noise. They have a tendency to "overfit" the data. The Bayesian method, which by its nature includes trade-offs between model complexity and fit to the data, avoids this problem. A Bayesian classification is the most probable set of classes given the data and the most probable assignment of each object to a class, and a program called Autoclass has been developed to implement this approach.

AUTOCLASS is based on a principle of inference first enunciated by Thomas Bayes in 1763. Suppose that  $D$  is a set of data and  $H_1, \dots, H_n$  are distinct hypotheses. Bayes Theorem says that the probability any one of the conditions, say  $H_k$ , occurs given  $D$  is the proportion of  $p(D)$  contributed by the  $k$ th term:

$$p(H_k | D) = \frac{p(H_k)p(D | H_k)}{p(D)}.$$

The probability of the data is obtained from the law of conditional probability,

$$p(D) = \sum_{i=1}^n p(D | H_i)p(H_i).$$

Bayes theorem is often stated in the form, "the posterior probability of the hypothesis given the data is proportional to the product of the prior probability of the hypothesis and the likelihood of the data given the hypothesis," where  $1/p(D)$  is the constant of proportionality. Bayes Theorem shows how to calculate a backward inference, sometimes called "reversed conditioning" or "inverse probability."

Now suppose that we interpret the  $H_i$  as possible models (hypotheses) that explain given experimental data  $D$ . Given any model, one can calculate the likelihood that the data will be observed in that model,  $p(D | H_i)$ . If one also has a value for the prior probability of each model,  $p(H_i)$ , one can then use Bayes Theorem to calculate the probability of each model given the data. It is then reasonable to say that the "best" model is the most probable one according to this calculation. This is called Bayesian inference.

In this approach, probabilities are interpreted not as frequencies observable through experiments, but as degrees of plausibility one assigns to each hypothesis based on the data and on one's assessment of the plausibility of the hypotheses prior to seeing the data.

The AUTOCLASS program applies Bayesian inference to determine the most probable classification of given data. In the case of the IRAS data, it assumes that the spectral intensity at each wavelength is accounted for by a normal distribution whose parameters come from one of  $N$  classes. Each class has a vector of parameters, the means and variances for each of the 94 intensities making up a spectrum. To specify a hypothesis, we need to state a value of  $N$ , a vector of 94 means and variances for each class, and a probability that each of the 5,425 records belongs to each particular class -- just over  $10^6$  numbers in all. Constructing a sample of several million hypotheses of this type from the astronomically large space of all possible hypotheses, and then applying Bayes Theorem to find the best, would far exceed the processing capacity of any existing supercomputer. Instead of an enumeration, AUTOCLASS uses a search procedure to iteratively modify a current hypothesis to obtain a maximum of Equation (1). The search procedure contains extra steps to attempt escape from local maxima. When it completes its search, AUTOCLASS has constructed a locally most likely hypothesis that explains the data. AUTOCLASS takes about 36 hours on a Symbolics computer to process the IRAS data.

As an important domain that illustrates the Bayesian approach, we developed the theory as applied to automatic classification of data. As a result of initial exploratory work, we implemented an automatic classification program (AutoClass II) that is sufficiently mature that it can be applied to many different databases. AUTOCLASS II can find classes with a combination of real-valued and discrete data and no prior information about what classes might be present. Unlike previous automatic classification programs, AUTOCLASS does not need to be told how many classes are present or even if there are any classes at all. AUTOCLASS uses a new extended Bayesian approach that searches for the most probable classification, and assures that any classes that are found have a real cause---i.e., classes are not an artifact of the search process.

## Discoveries in Astronomy

NASA's Space Science Division's Astrophysics Experiments Branch (Code SSA) has many largely unanalyzed databases which could benefit from this type of automatic classification. To test the theory, we applied AUTOCLASS to many real databases, with interesting results. The largest database analyzed was the IRAS Low Resolution Spectral database described above. The resulting classification recovered many well-known classes of stars as well as finding new classes of subtly different spectra of considerable astronomical interest. Independent information has since confirmed the validity of these new classes. This new classification was issued as a NASA reference publication #1217 in March 1989. As a result of these data analysis experiments, beginning in FY88 and continuing through FY89 and FY90, a number of limitations of the original model became apparent. In particular, the current AUTOCLASS assumes independence of the attributes within a class for simplicity, but this assumption is the major limitation in its performance. These experiments showed that the greatest improvement to AUTOCLASS would be allowing various models of depen-

dependency to be incorporated into the search. Beginning in FY89, we began work on methods for discovering significant correlations between variables, and building these correlations into the classification model.

The theory for discovery of significant correlations in data turned out to be much more complicated than was originally expected. However, we have made considerable progress in modeling the major forms of correlational dependency, and expect to publish a paper on the subject. This research is a first step to extending AUTOCLASS to handle attribute dependency.

In parallel with the development of dependency models, we have been rewriting the original AUTOCLASS II code to make it faster, more accurate and user-friendly. The process of turning experimental code into a robust usable system is a larger effort than the original code development. However, as the result of much work, a new version of AUTOCLASS (III) is now being integrated with the input interface, and should be available for general release by the end of April 1990. Other AUTOCLASS research is aimed at finding hierarchical classes, and integrating this capability with the discovery of inter-variable dependencies within an efficient search procedure.

In addition to extensions to AutoClass, the Bayesian theory can be extended in various ways that provide useful data analysis tools for NASA. In particular, we have begun a collaborative research project with code RFE aimed at discovering the basic turbulence structures (mainly vortices) in outputs from numerical fluid flow simulations. This work is initially funded by a DDF grant. Also, Dr. Wray Buntine has recently joined the group, to do research on the discovery of patterns that predict particular variables (supervised learning). This research is important because in many NASA applications, the goal is to make specific predictions rather than look for relationships among a set of variables, as in classification.

This work has been ongoing for four years. Members of the Bayesian Learning project team include: Peter Cheeseman (RIACS), John Stutz (NASA), Bob Kanefsky (Sterling Software), Robin Hanson (Sterling Software), Wray Buntine (RIACS).

# Networked Systems



View of a portion of Valles Marineris on Mars. Photographed from full color computer screen image generated by VPE project interactive 3D rendering software.



---

# Networked Systems

The charter of Networked Systems at RIACS is to assist NASA in the development of networked systems technologies that are critical to NASA missions in remote exploration and experimentation.

Future NASA missions such as the Mission to Planet Earth and missions to the moon and Mars involve teams of scientists, engineers, and operations personnel working with each other and with remote facilities. To support these missions adequately, new capabilities are required. For example, NASA will need to transmit massive amounts of data over high-bandwidth networks, new paradigms will be needed for interaction between ground and space-borne personnel, remote control of instruments will be necessary, and data visualization techniques will be needed for more rapid and effective data interpretation. RIACS works closely with NASA to help develop these capabilities.

## A. Overview

Research within Networked Systems has been directed towards distributed systems, collaboration technology, and workstation technology. In addition to investigating theoretical foundations of these technologies, we have conducted experimental validations of our results via multi-disciplinary pilot projects.

### *Distributed Systems*

Our research in this area has been directed towards developing an architecture for a system of distributed intelligent agents and towards analysis of high-performance networks, particularly those located at least partially in space.

Many systems for NASA's future missions will be composed of autonomous components embodying AI processes -- so-called "intelligent agents." A possible example might be the system of payloads on the Space Station; each payload draws on resources such as power and heating, and each can be delivered with a control agent containing knowledge of its function, demands, and safety requirements. On delivery, a payload control agent could identify itself to the network and become part of the local community of agents. Our work has focused on the architecture of such a system, especially the protocols for agent interactions.

The goals of the research are to:

Provide an execution environment that hides as much of the distributed environment from the programmer of agents as is feasible.

Simplify the programming of distributed agents by carrying over as many familiar programming concepts as possible into the world of distributed computation.

Provide a conceptually simple framework for programming the agents in those cases where no natural extensions to familiar programming concepts exist.

Networks provide the infrastructure for distributed systems. High-performance networks are vital to NASA, both to support its missions and to support its daily activities. Both low delay and high throughput are required for NASA's applications. Network analysis activities within the Networked Systems area have generally been directed towards analyzing performance of the Space Station Freedom Data Management System. We have also assisted in the development of recommendations for international data-system standards to support transmission of data over the space-to-ground link.

Networked Systems personnel working in this area are Marjory Johnson and Henry Sowizral.

### *Collaboration Technology*

NASA projects involve many groups of people at different locations working together. In the past, such projects have used traditional means -- telephones, travel to meetings, and, more recently, video-conferencing -- to maintain coordination, but they have generally not exploited modern technologies such as multimedia conferencing.

The use of computing technologies in NASA group operations is vital for future NASA missions. For example, scientists who fly experiments on Space Station Freedom want to access and manipulate their space-borne instruments from their laboratories on the ground. Telescience is a word that has been invented to describe this new and potentially effective mode of operation. Telescience also has potential for applications such as disaster management, to bring image data to the field manager during forest fires, earthquakes, or floods.

Our research in collaboration technology addresses the design of workflow among a group of dispersed scientists and the computational tools that support their collaboration. In addition to integrating existing technologies into the scientific work of selected NASA programs, projects, and missions, our group is developing new technologies and paradigms.

We have identified several opportunities to make emerging collaboration technologies available to the general NASA community. Our strategy has included:

Assisting the NASA Science Internet Project Office (NSIPO) at Ames to support the establishment of the National Research and Education Network (NREN) and the NASA Science Internet (NSI). The NREN, a network that will be fundamental to the support of remote interaction among people and resources, is being established by



Federal Agencies as part of the national initiative in High Performance Computing. NSI is intended to provide a comprehensive networking service for the OS-SA/science community through the use of government-sponsored and public networks.

Participating in networking committees, such as the Internet Activities Board (IAB), the Coordinating Committee for Intercontinental Research Networks (CCIRN), and the Federal Research Internet Coordinating Committee (FRICC).

Participating in the development of a National Collaboratory, to develop the infrastructure and procedures by which scientists at different locations can collaborate on research.

Collaborating with the Ames Telescience Project in the investigation of remote coaching in the context of Space Station Freedom. We are developing methodologies and evaluation strategies for ground-based scientists, mission control, and payload specialists to operate as a team conducting space experiments.

Networked Systems personnel working in this area are Maria Gallagher, Richard Haines, Vicki Johnson, and Barry Leiner.

### *Workstation Technology*

Workstations function as a door into a world of computational resources, agents, and people. The purpose of this project is to investigate workstation technologies to support NASA applications ranging from planetary visualization to Space Station Freedom payload operations.

The X window system is emerging as an important standard for workstations. By separating computing displays, it makes possible an environment in which machines that compute are different from machines that display. We are building a single environment for scientists based on X-windows, including the required tools for interactions between scientific databases, instruments, and other scientists. This environment is intended to support RIACS scientists as a beta site and will constitute the NASA telescience workstation environment. It builds on public software whenever possible.

Visualization is an important aid to understanding large data sets. We are exploring visualization of three-dimensional data in our investigations of Virtual Planetary Exploration, conducted jointly with NASA Ames Human Factors Branch (Code FLM). The goal is to use computer graphics systems coupled with novel input/output interaction systems, such as the NASA stereoscopic helmet display system, to allow the scientist to do "virtual exploration" by "roving through" or "flying over" visual simulations of terrain reconstructed with data collected from remote sensors. We are converting JPL and U.S. Geological Survey (USGS) planetary data for use on the project's graphics supercomputers (Stellar), designing function library routines for rendering and displaying 3D planetary scenes, and providing software documentation, and user's guides.

Networked Systems personnel working in this area are Robert Brown and Lew Hitchner.

### *Applied Pilot Projects*

Because of a tradition of "supply-side" thinking about innovations in computing, many technologies have never entered the domains of users. As a step toward "demand-side" thinking, we have applied user-oriented rapid-prototype testbedding to design computing technologies that support scientific users in their work. This methodology enables users and developers to design, build, and integrate a prototype system tailored to the user's requirements. The users then experiment with this prototype system. As observers we evaluate the effectiveness of the system and its use. In this way the users are able to articulate their own requirements precisely and the developers can adapt the software and systems to serve those needs in a more effective manner. As observers we are also able to surface new principles of design that can be used in other domains.

The Ames Telescience Program is a major example of "demand-side" design. It serves as a focal point for both Ames telescience-related activities and for the NASA/OSSA testbed activities in telescience. RIACS supports the Ames Telescience Project Office in the planning and coordination of the Ames Telescience Project and OSSA Telescience Testbed Project.

We also collaborate with the earth sciences community both at Ames and nationally, evaluating technologies that can support future activities in global environmental research. This effort will involve many organizations and computing systems worldwide. We have also undertaken a pilot study to use automatic servers to help researchers in the parallel-processing area to share their results.

Networked Systems personnel who have contributed to this area include Robert Brown, Dee Doyle, Maria Gallagher, Richard Haines, Vicki Johnson, Richard Johnson, Barry Leiner, Henry Sowizral, and James Woods. Other members of RIACS who have participated in these projects are Anne Kohutanycz and Michael Raugh.

## **B. NASA Sponsors and Collaborators**

The projects we have undertaken would not be possible without the support and cooperation of a wide variety of people at Ames, elsewhere in NASA, and throughout the scientific and technical community.

In particular, we would like to acknowledge our Ames sponsors and collaborators:

Computer Systems and Research Division (Code RC)

*Marcelline Smith*

Information Sciences Division (Code RI- now Code FI)

*Henry Lum, Jr., Alan Fernquist, Terry Grant, Ellen Ochoa, Jerry Yan*

Advanced Space Technology Office (Code DS)

*William Berry, Tony Gross, David Morse, Daryl Rasmussen*

Earth System Science Division (Code SG)

*James Lawless, David Peterson, Jim Brass*

Life Science Division (Code SL)

*Edwin Force*

Information and Communications Systems Division (Code ED)

*Dale Lumb, James Hart, William Jones, Fred Rounds*

Aerospace Human Factors Research Division (Code FL)

*Michael McGreevy*

We would also like to acknowledge our supporters in NASA Headquarters and other organizations:

Information Systems Division (HQ/RC)

*Lee Holcomb, Paul Smith*

Communications and Information Systems Division (HQ/EC)

*Ray Arnold, Joe Bredekamp, Tony Villasenor*

Earth Sciences Division (HQ/EE)

*Shelby Tilford*

Space Station Program Office

*Mel DeGree*

Space Station Level I

*Carolyn Griner, Stan Fishkind, Mike Pasciuto, Gregg Swietek*

## C. Accomplishments

### *Distributed Systems*

- **Distributed Intelligent Agents**

Henry Sowizral

During the past eighteen months, Sowizral studied various distributed operating systems and identified gaps in current development tools that complicate the construction of agents. He developed a prototype using one element of the overall architecture to guarantee the order of message delivery among agents.

In the future Sowizral plans to develop a more sophisticated prototype system of agents, consisting of both machines and people. This system will require the implementation of distributed ordering, but allows delaying the implementation of distributed continuations and self-reflection. He will document the behavioral patterns of the human participants in the prototype system. He will also demonstrate a preliminary version of distributed continuation syntax and semantics. He will investigate the on-board processing requirements for such systems.

The Information Sciences Division at NASA Ames sponsored this activity.

- **Network Analysis**

Marjory Johnson

M. Johnson developed and documented an argument that NASA needs to conduct networking research to ensure that unique networking requirements of its space missions can be satisfied. Space Station Freedom related analyses include an analysis of the Operations Management System, a presentation of modeling activities supporting development of the Data Management System, and an analysis of data that will be generated by the Space Station Freedom System. The results from these analyses will be published in the literature and presented at various conferences.

In addition to her Space Station related activities, M. Johnson presented an invited tutorial, "Network Protocol Performance," at the ACM SIGMETRICS '89 conference. She was invited to present an updated version of the same tutorial at the ACM SIGMETRICS '90 conference.

M. Johnson is chairperson for an IFIP-sponsored Second International Workshop on Protocols for High-Speed Networks, to be held in Palo Alto, CA, on November 27-29, 1990. Forty research papers have been submitted for this workshop. M. Johnson, with support from the workshop program committee, is currently planning the program for the workshop.

Next year M. Johnson will continue her work with the NASA Ames Intelligent Systems Branch of the Information Sciences Division to develop strategies for assuring timely network response for fault management of the Space Station Freedom Data Management System and to participate in other Advanced Architectures testbed activities.

Sponsors for these activities include the Information Sciences Division and the Computer Systems and Research Division at NASA Ames, the Operations and Utilization Division at Space Station Level I, and DARPA.

- **CCSDS Activities**  
Marjory Johnson

The Consultative Committee for Space Data Systems (CCSDS) is an international organization that develops recommendations for data-system standards to support space missions. Member agencies include NASA/USA, the European Space Agency (ESA)/Europe, and space agencies in the United Kingdom, Canada, France, West Germany, India, Brazil, and Japan. M. Johnson's role in the group is to provide analytical expertise to support the development of these recommendations. Recent activity within CCSDS culminated in the development of data-system standards for Advanced Orbiting Systems (AOS), such as Space Station Freedom.

CCSDS Path Service was developed to streamline the transmission of large volumes of telemetry data from space to ground. Working with other CCSDS members, M. Johnson conducted a performance analysis of this service, comparing it with general-purpose networking protocols. The analysis showed that Path Service possesses significant advantages in terms of protocol-processing requirements and channel utilization. The results of this analysis were published in a paper entitled "Performance Analysis of CCSDS Path Service."

Because of the complexity of the CCSDS Recommendation for Advanced Orbiting Systems, several supporting documents accompany the main recommendation. M. Johnson has served as editor for two of these documents: "CCSDS Report Concerning Space Data System Standards, Advanced Orbiting Systems, Networks and Data Links: Summary of Concept, Rationale, and Performance" (the AOS Green Book), and "CCSDS Recommendation for Space Data Systems Standards. Advanced Orbiting Systems, Networks and Data Links: Cross Support of Management and Signalling."

In April, 1989, M. Johnson hosted an international meeting of CCSDS Panel 1 at Ames. This was primarily a working meeting to finalize the Recommendation for

Advanced Orbiting Systems. An audio/video workshop was held in conjunction with the Panel 1 meeting.

Next year M. Johnson will continue as editor for the CCSDS AOS Green Book, and she will participate in planning for future activities within CCSDS. Her plans also include experimental validation of the CCSDS Advanced Orbiting Systems protocols via telescience experiments.

These activities have been sponsored by the Space Station Program Office and the Advanced Space Technology Office at NASA Ames.

- **Space Station Freedom Testbed Activities**

Marjory Johnson, Henry Sowizral

We have recently initiated a project to analyze performance of the Space Station Freedom end-to-end communications system, a networking system that includes space, ground, and space-to-ground components. This analysis will include performance of hardware and software components of the on-board Data Management System. We plan to incorporate our work into the Advanced Architectures testbed activities within the Information Sciences Division at Ames.

These activities are sponsored by the Operations and Utilization Division at Space Station Level I and the Advanced Space Technology Office at NASA Ames.

## *Collaboration Technology*

- **NASA Science Internet Project**

Barry Leiner, Maria Gallagher

Leiner has assisted in the planning for the National Research and Education Network. He participated in the IAB and the CCIRN and served as the IAB observer to the FRICC. He supports ongoing strategic review of the NSI architectural and organizational issues, determining best approaches for supporting the networking requirements of the science community. He completed a study on the policy issues involved in interconnecting agency networks. The results from this study, which included two workshops held in 1988, were published as RIACS report [TR89.25]. Under subcontract SRI International performed a supporting study of technical, organizational, and cost-management issues for science networking. These results were documented in presentations and reports to OSSA management and were factored into planning for NSI management.

Leiner participated in an NSF-sponsored workshop that formulated the research agenda for the National Collaboratory. NASA's Mission to Planet Earth and the associated Global Environmental Change research program were identified as major initial testbeds for the Collaboratory. He played a major role in documenting the results of this workshop. The resulting report, "Towards a National Collaboratory," is available upon request. A more detailed report on the interactions with scientific instruments is available as RIACS TR89.45.

Leiner and Gallagher have supported both the NREN and the NSI by working with the NASA Science Internet Project Office at Ames. Gallagher served as liaison between the NASA Science Internet (NSI) and telescience communities, assisting

NASA telescience participants in defining and representing their requirements to the NSIPO. She also undertook preparation of a primer for the use of NSI by the science community.

In the future, we will continue to support strategic planning for the introduction of new technologies into NSI. We will continue participation in research networking coordination with the IAB and CCIRN. We will complete the NSI User Primer and assist in the preparation of other user-oriented documentation. We will continue acting as the NSIPO interface to the telescience community for defining networking requirements and inserting them into the NSI engineering process.

The sponsor for these activities is the Information and Communications Systems Division at NASA Ames.

- **Remote Coaching Facility**

Richard Haines, Vicki Johnson

Haines conducted a prototype experiment simulating data, voice, and video links among the three groups (on-board crew, ground mission control, and ground-based scientists) performing experiments in life sciences. In this way we gained considerable insight into the design of experimental protocols where the manipulator (Mission Specialist) is not an expert in life sciences and must be coached by principal investigators from remote locations. Vicki Johnson developed a "distributed checklist" facility on Macintosh computers to assist these groups in carrying out their experiments in the face of unreliable communication channels. We believe this to be a common scenario for certain telescience applications, and therefore, as part of the workstation technology research activities (see below), Brown and Doyle built a prototype "distributed checklist" tool using X-Windows that permitted all participants to follow the progress of the experiment.

Haines evaluated the Remote Coaching facility described above, through a systematic experiment that introduced failures in audio, video, and data links. He formulated a benchmark, useful in future studies of this kind, that incorporates human and system parameters that have not been previously considered. The work in remote coaching experimentation was documented through technical reports (RIACS TR89.01 and TR89.31) as well as conference papers and presentations primarily by Haines.

Haines (with NASA engineer Robert Jackson) conducted a controlled experiment to quantify video bit-rate requirements to support remote small-animal monitoring on Space Station Freedom. Valuable insights were gained and documented in RIACS TR90.19. The Advanced Space Technology Office at NASA Ames is the sponsor for these activities.

## *Workstation Technology*

- **Workstation Development**

Robert Brown, Dee Doyle

Brown and Doyle have continued investigations into open software workstation environments. We evaluated a variety of software packages for telescience applications;

specifically the X Window system from MIT, the Motif toolkit from the Open Software Foundation, the Transportable Application Executive (TAE-Plus) from Goddard Space Flight Center, and assorted user interface toolkits for the X Window System. We brought the X Window system into RIACS and have offered it for general use throughout NASA Ames. In support of the Telescience Testbed Program, we provided consulting services in workstation development tools for the testbed participants. An initial telescience workstation environment based on SunView is documented in RIACS TR89.01. This environment was upgraded in cooperation with RIACS facility staff to an X Windows-based system. Brown served in the elected position of Vice-Chairman of the Sun Users Group. In the future Brown and Doyle plan to produce a paper entitled "Paradigms and Packages for Telescience Workstation Applications" that can be used by experimenters to make software-design decisions for their telescience projects. We will produce several more prototype applications that can be used as templates for others' work. The sponsor for these activities has been the Advanced Space Technology Office at NASA Ames.

- **Workstation Workshop**

Barry Leiner, Robert Brown, Dee Doyle

In March of 1990, RIACS hosted the Workshop on NASA Workstation Technology, which was attended by approximately 100 people from around the agency. The purpose of the workshop was to bring together those people within NASA who are active in the development and application of advanced workstation technology, in order to share ideas and experiences. There were two days of plenary sessions and a third day of working meetings to create a report about the state of workstation technology within the agency. On the evening of the first day, we hosted a vendor and demonstration session, where six workstation vendors displayed their products and several NASA-sponsored software systems were demonstrated. Brown designed and built the technical program for the workshop, Doyle was responsible for overall organization and logistics, and Lorraine Fisher managed local arrangements.

This workshop was sponsored by the Advanced Space Technology Office at NASA Ames, along with RIACS core funds.

- **Virtual Planetary Exploration Project**

Robert Brown, Lew Hitchner

Brown and Hitchner actively participated in the Ames Visualization for Planetary Exploration (VPE) project. We completed the phase one analysis and synthesis computer software for modification and management of the digital terrain data. We used Mars Digital Terrain Model (DTM) elevation data and Mars Digital Image Model (DIM) surface reflectance image data that was processed by the USGS (Flagstaff) and obtained from the National Space Science Data Center (NSSDC) at NASA Goddard. This work was performed using two Stellar (now named Stardent) GS1000 graphics super-computers. During the year one of these was upgraded to a GS2000 dual channel system (for eventual use to support stereoscopic image generation). We also developed the software interface for two interactive input/output devices -- the Polhemus six degree-of-freedom position tracker (input) and the Ames head mounted display (output). A software application was developed that offers

users a 2D plan view of selected sub-regions of planetary terrain around which a scientist can interactively roam at selected latitude and longitudes. Another software application was developed that shows a user a 3D view of a selected terrain scene and permits a scientist to interactively travel around in the virtual landscape. Work was begun on scene synthesis software for a wide field-of-view 2D static image called a "cylindrical frame buffer." All our development has been done in a networked workstation environment using industry standard X11 windows software plus locally written code. We also developed hardcopy photographic output capability for producing still images.

Future work on the VPE project will include further development of the software for 2D plan views, the 3D interactive simulator, and the "cylindrical frame buffer." A major near term effort is design and development of the "exploration user interface" that will make our system intuitive to use by planetary scientists. Upgrading our software to include stereoscopic display capability will be completed. We will continue our development of a database management system for multi-resolution organization and storage of terrain data. We will design and implement a terrain data storage and retrieval system incorporating hierarchical storage organization and data compression and decompression. We will enhance the performance of our 3D simulator software by adding rendering enhancements that utilize knowledge of the terrain data structure to provide faster animation rates. We will investigate remote display of VPE terrain scenes via networked graphics workstations (tele-visualization). During the Magellan orbiter radar mapping mission we will obtain Venusian terrain data from the JPL mission team and display it for their visualization on a VPE workstation.

This project is sponsored by the Aerospace Human Factors Division at NASA Ames.

### *Applied Pilot Projects*

- **Telescience Coordination and Planning**  
Barry Leiner, Maria Gallagher

During 1989 Leiner and Gallagher completed the Telescience Testbed Pilot Program (TTPP). This program brought together 15 universities and 8 NASA centers to investigate telescience requirements as these researchers conducted prototype experiments over networks. We acted as observers of the scientific principles needed for the design of information systems to support effective telescience. We issued a three-volume final report integrating the results from the program (RIACS TR89.07, TR89.08, and TR89.09). In addition, the participating organizations collectively published 62 technical reports and papers.

Leiner and Gallagher supported the Ames Telescience Program Office in the planning and conduct of the Telescience Testbed Project (the follow-on to the TTPP) and in the general advocacy of telescience throughout NASA and the university community. Gallagher keeps the Office of Space Science and Applications (OSSA) community informed about telescience developments by contributing news items regularly to the OSSA Information Systems Newsletter.



We hosted a meeting of the participants in the TTP project in the fall of 1989. Gallagher articulated the TTP requirements for networking support to the NSI Project Office (see above). Finally, we assisted in the development of an overall Ames approach to telescience, documented in the Ames Telescience Program Plan.

This project has been sponsored by the Advanced Space Technology Office at NASA Ames.

- **Remote Coaching Facility**

Richard Haines, Vicki Johnson, Dee Doyle

Working with members of the Ames telescience project and the Life Sciences Division, Haines, V. Johnson and Doyle developed an experimental facility to test alternatives for carrying out experiments requiring remote coaching. This has resulted in concrete results (see above under collaboration technology) in the requirements for data, video, and audio in support of remote interactions and has demonstrated the viability of such an approach to the life sciences and Space Station Freedom programs. Haines conducted human factors studies of multi-media tools in this system. He evaluated the effectiveness of linked computer workstations, two-way video, and open-mike audio systems on productivity, work-load, and error-rate reduction. This life sciences remote-coaching facility was demonstrated to NASA's 1.8 m centrifuge facility system study review committee, work package 1 telecommunications planning group, Marshall Space Flight Center and Johnson Space Center communication and tracking personnel, JPL robotics lab personnel, Lockheed Missiles and Space Corporation personnel, Boeing space group personnel, McDonnell Douglas Corporation (Huntsville), Canadian Astronautics Ltd., General Electric (JSC), and several dozen universities. We are working with Ames Life Sciences personnel to further identify Space Station Freedom bandwidth requirements to support remote coaching operations.

This activity is sponsored by the Advanced Space Technology Office at NASA Ames.

- **Global Environmental Change Project**

Richard Johnson, Barry Leiner, Robert Brown

The proposed Mission to Planet Earth requires effective methods for people at many institutes to share resources, couple models, and collaborate. R. Johnson, working with the Ames Earth Sciences Division, put together a plan for a consortium to promote this. He is now planning a pilot project to be conducted jointly with the Ecosystems Division to address this area.

As General Chairman, R. Johnson planned and organized the American Astronautical Society Meeting in March 1989 entitled "Global Environmental Change: The Role of Space in Understanding Earth." As Vice-chairman of the FCCSET Committee on Earth Sciences, R. Johnson contributed to the planning and writing of the report, "Our Changing Planet: The FY1990 Research Plan," July 1989, and to a similar report in executive summary form, "Our Changing Planet: The FY 1991 Research Plan," January 1990. The plans form the basis for the U.S. global change research programs.

Leiner, R. Johnson, and Brown participated in the Ames planning for an Ames Center strategy and program initiatives for the Global Change Technology Initiative (GCTI). R. Johnson and Brown attended the NASA-wide GCTI Data Systems Workshop, Bethesda, Maryland, May 1989.

In the future we will collaborate with NASA Ames management and researchers to investigate the utility and feasibility of conducting atmospheric and ecosystems research with long-duration, high-altitude aircraft and to investigate options for utilizing existing and planned facilities, data bases, personnel, and programs as the foundation for an integrated pilot study of bio-geochemical systems dynamics. In another area we are working with Ames Earth System Science Division on a project to introduce collaboration tools into disaster management. The potential roles of advanced computer technologies, high-bandwidth communications networks, and inter-institutional relationships for global environmental research will also be investigated. This activity has been sponsored by the Earth System Science Division at NASA Ames.

- **User Support at Supercomputer Facilities**

Henry Sowizral

Sowizral examined facilities for user support at national supercomputer sites because we see that automated help and consulting can help NASA's growing and diverse community of supercomputing scientists. We found that a single organization serving multiple supercomputer sites could significantly reduce overhead costs and more efficiently utilize highly trained staff. In this initial study of distributed user services, Sowizral visited three sites, interviewed the personnel providing user services, interviewed users, and examined the software and literature they provided for their supercomputer users. His report identified key technologies that would enable the development of a distributed user-services organization. He recommended a pilot program. This activity was supported by the Numerical Aerodynamic Simulation Systems Division at NASA Ames.

- **Unified Information System Project**

Barry Leiner, Robert Brown, Dee Doyle, Anne Kohutanycz, James Woods

With support from DARPA, we will develop a unified information system for the parallel processing research community. It will enable them to share and discuss research findings rapidly. We will begin with the Connection Machine research community. We will survey them to determine what information must be in the repository and what operations are needed to add, modify, delete, or discuss that information. We will design a naming scheme that will allow users to access information stored in associated repositories without encoding location dependencies.

This work is sponsored by DARPA and the Numerical Aerodynamic Simulation Systems Division at NASA Ames.

- **Parallel Processing Database**

Dee Doyle, Mike Raugh, Anne Kohutanycz

This is a new activity, whose purpose is to develop a database of available parallel processing software and algorithms which can be used as a tool in planning for the High Performance Computing Initiative.

We have made the following progress: Sybase has been selected for use as the database, Kohutanycz has attended a class on Sybase database administration and is in the early stages of database design, and Raugh is contacting members of the parallel processing community to identify important algorithms.


In the future we plan to expand the database to contain entries for both public domain and proprietary software. Software developed at NASA centers, universities, government laboratories, and industry will all be included.

This activity is sponsored by the Office of Aeronautics and Exploration Technology at NASA Headquarters.



# Parallel Systems

---



Close up of unstructured grid about airfoil with extended flaps generated by Dennis Jespersen of NASA Ames.



---

# Parallel Systems

The charter of the Parallel Systems of RIACS is to engage in fundamental studies of the algorithms and software needed to solve large-scale NASA scientific problems on available and projected parallel supercomputer architectures. These studies are expected to make significant contributions in the state-of-the-art techniques in computational fluid dynamics including 3-D Navier Stokes unsteady flow and incompressible high altitude transonic flow, in computational chemistry, and other disciplines.

In pursuing our charter, we strive to become a recognized leader in selected areas of research critical to NASA's priorities by recruiting and supporting top scientists, by maintaining high standards for appointment and promotion, by participating in meetings, symposia, workshops and other professional activities, and by developing strong relationships with NASA, universities, government, and other organizations.

During 1989, our effort was allocated among four separately-funded projects. The Advanced Algorithms and Architectures (AAA) Project is supported by the NAS Systems Division of NASA Ames Research Center. The objective of the AAA Project is to design and test new algorithms that bring forth the full computational power of massively parallel architectures, to extend the theory and concepts of fast algorithms for parallel machines, to analyze the performance and usability of new systems including the mapping of specific problems onto target architectures, and to identify the strengths and weaknesses of new architectures.

The Center for Advanced Architectures (CAA) Project is supported by DARPA and administered by the NAS Systems Division of NASA Ames Research Center. The objective of the CAA Project is to conduct research into the applicability of novel computer architectures (e.g., the Connection Machine and Intel Touchstone hypercube) for the solution of computationally intensive problems arising in a variety of scientific disciplines. Like the AAA Project, it is interested in efficient algorithms for massively parallel computer architectures; but unlike the AAA Project, it focuses on these questions for specific machines of joint interest to NASA and DARPA.

The Parallel FFT and Matrix Algorithms Project is supported by the NAS Systems Division of NASA Ames Research Center. The objective of this project is to develop efficient variations of the FFT on the Connection Machine and other parallel computers and to examine efficient parallel algorithms for computing eigenvalues of tridiagonal matrices.

Finally, the RIACS Core Project was supported by a portion of the RIACS core funds. The objective is to initiate new and innovative research projects especially collaborative efforts and technology transfer mechanisms.

A large number of scientists, visitors and students participated in these four projects. The AAA project involved Paul Frederickson, Youcef Saad, Robert Schreiber and Richard Sincovec of RIACS along with visitors: Niel Madsen, Andy Wathen, and Tony Chan. The CAA project had a large number of participants. From RIACS there were Robert Brown, David Curry, Roland Freund, Steve Hammond, Chris Lewis, David Rogers, Robert Schreiber, Richard Sincovec, and Carl Williams. Visitors and students were: George Adams, Leo Dagum, David Day, Jack Dennis, Billy Stewart, Walter Tichy, Charles Tong, Ray Tuminaro and Harry Wijshoff. RIACS participants in the Core Project included David Curry, Roland Freund, Robert Schreiber and Richard Sincovec of RIACS and student Bill Behrman. Paul Swarztrauber was the principal investigator for the Parallel FFT and Matrix Algorithms Project.

## A. Objectives

Parallel Systems has five major objectives and a number of strategies for achieving the stated objectives.

### *Objective One: Parallel Linear Algebra Algorithms*

Conduct research and development in parallel linear algebra algorithms with emphasis on the solution of the equations arising in CFD calculations.

This is the current major effort and unifying theme of the Parallel Systems. Improving linear algebra algorithms for parallel machines will make a major contribution since the most computationally intensive part of many scientific and engineering problems involves solving a linear algebra problem. In CFD calculations, the linear algebra problem results from the discretization of the partial differential equations that describe the fluid flow. Often this linear algebra problem has a sparse structure that one takes advantage of in the algorithms and their implementation on parallel computers. The combination of the expected hardware speedups plus these algorithm improvements makes the goal of a 1000-fold speedup (relative to today's Cray 2) for CFD calculations by 1995 reachable. Such a speedup is essential for NASA to attain one of its grand challenges--flying an airplane inside a computer.

### *Strategies for Achieving Objective One*

Develop efficient iterative methods for general nonsymmetric linear systems by focusing on preconditioning techniques.

Develop efficient implementations of sparse matrix solvers on massively parallel computer architecture.

Adapt robust classic multigrid solvers to the needs of a fully implicit or continuation method based flow solver.



### *Objective Two: Parallel Algorithms*

Develop new parallel numerical algorithms for solving problems of interest to NASA on parallel computers.

Parallel machines with thousands of processors present new challenges to algorithmic designers. Large problems must be decomposed into components that can be worked on by separate processors, and the programming environments must keep track of those components. Data flow and communication among the program components must be minimized to achieve performance commensurate with the hardware power. The computational load must be balanced among the processors, not an easy problem when the distribution of computational requirements among program components is data dependent. These challenges and others must be addressed in the design of parallel algorithms that deliver most of the hardware power.

### *Strategies for Achieving Objective Two*

Develop efficient Fast Fourier Transform (FFT) algorithms for highly parallel computers with applicability to important problems in CFD, signal processing and other areas.

Explore and evaluate techniques for solving partial differential equations on parallel computers; important techniques include grid refinement, independent time steps, and structured versus unstructured grids.

Design and implement new parallel algorithms for computing the eigenvalues of matrices.

Conduct research on particle simulation methods on parallel computers for high Mach number, high altitude, rarefied gas flow.

Model the convective flow of the earth's mantle in three dimensions using an octahedral triangulation.

Develop a collaborative effort with appropriate NASA scientists on numerical problems in control theory.

### *Objective Three: Architectures*

Evaluate new and important parallel computer architectures for use in solving problems of interest to NASA.

Currently there are many vendors marketing several significantly different computer architectures. It is important to NASA to identify those architectures that hold the most promise for achieving the performance required for NASA to solve the computational grand challenge problems that have been identified.

### *Strategies for Achieving Objective Three*

Participate with NASA in assessing the current and projected state of Cray-class supercomputers and parallel research computers including software.

Explore the use and performance of new parallel architecture and algorithms.

Examine innovative solutions of the Navier-Stokes equations using parallel computer with thousands of processors.

#### *Objective Four: Usability and Environments*

Conduct research in basic and applied computer science that enhances the usability of parallel computer architectures.

Software development environments, languages, compilers, debuggers, and tools in general for parallel computers are often inadequate. It is important to evaluate existing software and propose new software so that scientists and engineers can more effectively use parallel computers.

#### *Strategies for Achieving Objective Four*

Explore the impact of language on the development of parallel software.

Complete the Paradigm (SISAL) compiler for translating SISAL dataflow programs into machine code for the Connection Machine, and demonstrate its power for scientific applications of interest to NASA by testing it on the Large Eddy Simulation (LES) code (rewritten in SISAL).

Define communication primitives that can be assembled under algorithmic control to maximize internal communication bandwidth for a wide class of common communication tasks including matrix transpositions, perfect shuffles, and bit reversals.

Develop effective graph embedding heuristics and algorithms to map realistic problems, especially those arising from fluid flow codes with unstructured grids, onto parallel computer architectures.

#### *Objective Five: Technology Transfer*

Provide technology transfer. By "technology transfer" we mean the integration of computer science and technology expertise into NASA missions.

#### *Strategies for Achieving Objective Five*

Every member of the Parallel Systems will identify at least one NASA counterpart to work and collaborate with on problems of mutual interest.

Report research results in writing by technical reports and publications and by presentations at meetings, workshops, short courses, and seminars.

Develop a visitor program including summer students and faculty, consultants, student interns, faculty sabbaticals, and industrial affiliates.

Develop respected collaborative research programs with external organizations and with NASA constituents that are the foundations of RIACS's programs.

Sponsor and organize the RIACS/NAS seminar series and participate in the organization of meetings and workshops.

Make available software resulting from the research effort.

These strategies are all conducted against a background of five fundamental questions:

Which architecture provides the best performance for a given class of problems?

How can software technology be used to exploit the power of the architecture?

How can problems be automatically mapped to concurrent architectures?

How can the efficiency and usability of architectures be evaluated?

How should languages and tools be designed for effective use of new architectures?

## B. Accomplishments

### *Parallel Linear Algebra Algorithms*

- **Efficient Preconditioned Conjugate Gradient on Parallel Computers**

Steve Hammond, Robert Schreiber

The focus here is on sparse triangular systems. First, we analyze the dependence graph and reorder the index set so that independent tasks can be executed in parallel. Second, we dynamically schedule tasks for better load balancing and improved parallel efficiency. We explored different approaches for exploiting parallelism in the ICCG (Incomplete Cholesky Conjugate Gradient) method for solving large sparse symmetric positive definite systems of equations on a shared memory parallel computer. We focused on solving the sparse triangular systems that arise when using a preconditioner based on an approximate factorization of the coefficient matrix. An important difficulty in solving general sparse triangular systems is that the available parallelism depends on the zero structure of the matrix, and is therefore not known at compile time. The concurrency is data dependent and can be determined only at run time.

We showed that a straightforward analysis data dependencies allows scheduling that drastically improves the parallel efficiency. The idea is as follows. We permute (reorder) the index set of the recurrence equation for the triangular solve and put the indices in a queue. The processors repeatedly take indices from the queue, perform the associated calculations, and then take another index until all unknowns have been computed. Data dependencies are resolved by semaphores. We check indices in a shared array that indicate whether each of the unknowns has been computed. If a calculation depends on data that has not been computed, then the processor performing the calculation must enter busy wait in the shared array.

We showed that there is a tradeoff between storing the lower triangular part of a symmetric matrix and storing the entire matrix. Storing the lower part to save storage complicates the multiplication since both outer products (which require synchro-

nization) and inner products must be performed. The synchronization overhead slows down this operation. Again, we found that a straightforward analysis for data dependencies can lead to high parallel efficiency. Moreover, when the sparsity structure of a triangular matrix is irregular, a reordered dynamic scheduling method performs more efficiently than a reordered static scheduling method. [RIACS TR89.24.]

- **Sparse Solution Algorithms for CFD Applications**

Paul Frederickson, Bill Behrman

The objective is to develop a sparse implicit solver adequate for the Newton step in the continuation method on a CFD application. Continuation methods are essential for the study of multiple solutions of the steady Navier Stokes equations beyond a bifurcation point where a single nonlinear solution splits into more than one solution. Also continuation methods can be expected to provide a competitive numerical technique in flight domains where the usual algorithms converge with difficulty. The first phase of this study, the construction of a testbed for the analysis of the variational equations and the study of possible solution algorithms, was completed last summer with the assistance of Stanford graduate student Bill Behrman. A new multi-level smoother is under development which we expect will provide fast convergence for the strongly non-symmetric problem we have taken on as a test problem.

- **Study of Structured and Partitioned Indefinite Hermitian Matrices and the Development of Preconditioners for such Matrices**

Roland Freund, Thomas Huckle/Wuerzburg

Structured and partitioned indefinite Hermitian matrices come up in important applications such as nonlinear programming or in connection with the numerical computation of incompressible fluid flow. Typically the resulting matrices are partitioned into four blocks where the block in the lower right corner is zero. We proposed the new concept of quasi-spectral decompositions which is the natural transformation for partitioned Hermitian matrices of this type. On the one hand, this decomposition provides a tool for the study of connections between the inertias or eigenvalues of the matrix and its principal submatrices. In particular, using the quasi-spectral decompositions, one obtains very simple proofs of all known results involving connections of this type, as well as, a few new results. On the other hand, via incomplete quasi-spectral decompositions, it can be used to compute preconditioners for partitioned indefinite Hermitian matrices. [RIACS TR89.55.]

- **Multigrid Algorithms on Parallel Processing Systems**

Ray Tuminaro/Stanford, Tony Chan/UCLA

Ray Tuminaro developed a new multigrid-like method for the solution of the two-dimensional steady Euler equations describing flow around an airfoil. The method is similar to a standard FAS multigrid scheme in that convergence on the finest grid is accelerated by iterations on coarser grids. In this new method, however, a fine grid subproblem is processed concurrently with the coarse grid subproblem to further accelerate convergence. The new approach has the advantage that greater parallelism is possible than within a standard multigrid method making the algorithm more suit-

able for massively parallel machines. As a test case, we implemented Jameson's well-known FLO52 algorithm. Numerical experiments illustrate the new method's fast convergence. A detailed evaluation of the new method was performed based on numerical results and a mathematical execution time model. [RIACS TR89.41.]

- **Broyden Type Methods for the Solution of Linear and Nonlinear Systems of Equations**

Roland Freund, Peter Deufhard and Artur Walter/Konrad-Zuse Institute, West Berlin

It is well known that, for Hermitian positive definite matrices, most quasi-Newton methods are closely related to the conjugate gradient algorithm. The goal of this project is to investigate similar connections between quasi-Newton methods and conjugate gradient type algorithms for general non-Hermitian linear systems or nonlinear systems. So far, we have focused on Broyden's method. We have shown that two subsequent steps of Broyden's method can be viewed as a combination of one step of Richardson iteration and one step of conjugate gradient iteration with changing preconditioning matrix. This insight led us to natural generalizations of Broyden's method which combine multiple conjugate gradient or minimum residual steps with quasi-Newton techniques which improve the preconditioning matrix in the course of the iteration. First numerical test with such schemes have proven to be competitive with conjugate gradient approaches. It is planned to extend these algorithms to nonlinear systems and explore their use for the solution of the equations arising in CFD.

- **Multigrid as a Fast Poisson Solver**

Paul Frederickson

A clean Fortran version of the classic multigrid solver FAPIN (for Fast Approximate Inverse) was demonstrated to run at over two gigaflops on Reynolds, the Cray YMP 832 at NAS. With only  $75n$  floating point operations per iteration where  $n$  is the size of the system, this provides a faster solution, when the initial guess is sufficiently good, than an FFT based fast Poisson solver. The solution requires about four times as many floating point operations per iteration as evaluating one residual.

- **Approximation Problems which Arise in Connection with Polynomial Based Methods for Solving Linear Systems**

Roland Freund, Bernd Fischer/Stanford

Many iterative schemes for solving large sparse linear systems  $Ax=b$  are polynomial based methods: they produce approximate solutions to  $Ax=b$  which are of the form  $p(A)r$  where  $p$  is a polynomial and  $r$  some starting vector. For example, conjugate gradient type algorithms and semi-iterative schemes belong to this category. The analysis and the optimal design of polynomial based methods is intimately connected with the study of certain approximation problems for polynomials. In particular if  $A$  is non-Hermitian, polynomial approximation problems on sets in the complex plane result. In general, very little is known about problems of this type. Here, we studied these polynomial approximation problems for the important special case of ellipses in the complex plane. We obtained the complete solution of these problems for

"almost all" cases. Moreover, somewhat surprisingly, we proved that a widely referenced and supposedly well-known result that suitably normalized Chebyshev polynomials are optimal for approximation problems of this type is wrong in general. Finally, as a by-product, a number of new inequalities for complex polynomials were obtained. [RIACS TR89.04, TR89.05, TR89.17, TR89.21.]

- **Efficient Conjugate Gradient Type Methods for Solving Linear Systems with Complex Coefficient Matrices**

Roland Freund, Steve Hammond

Large sparse linear systems with complex coefficient matrices arise in important applications, such as electromagnetics, inverse scattering, numerical solution of time-dependent Schroedinger equations, and quantum chromodynamics (QCD). Typically, the matrices which arise in these applications exhibit special structures. For instance, they are complex symmetric or are shifted Hermitian matrices. However, these structures are usually ignored and the matrices are treated as general non-Hermitian matrices. In contrast to this standard approach, we develop and investigate conjugate gradient type schemes which exploit the special structures of these complex linear systems.

So far, our main focus was on basic conjugate gradient type algorithms for shifted Hermitian matrices and complex symmetric matrices. We devised and analyzed schemes which are based on a variant of the Lanczos recursion for complex symmetric matrices. We proposed a new approach with iterates defined by a quasi minimum residual property.

Most of the standard preconditioning techniques do not preserve the special structure of complex linear systems. More work needs to be done to design efficient structure-preserving preconditioners. Also, in the future, we will focus on iterative schemes for the special class of matrices which arise in QCD. [RIACS TR89.54.]

- **Efficient Polynomial Preconditioners for Hermitian and General Non-Hermitian Matrices**

Roland Freund, Steve Hammond, Thomas Huckle/Wuerzburg

The main challenge in designing efficient conjugate gradient type algorithms for parallel computers is the search for preconditioners which are suitable for parallel and vector architectures. Polynomial preconditioning is a technique that attempts to approximate the inverse of the coefficient matrix by a polynomial in the coefficient matrix. The coefficient matrix is not modified as required in incomplete factorization techniques. This makes polynomial preconditioning a very natural approach for massively parallel machines such as the Connection Machine. The fundamental problem that arises in this context is how to choose the polynomial preconditioner in order to optimally speed up the basic conjugate gradient type iteration.

The goal of this project is the development of such optimal polynomial preconditioners for Hermitian and general non-Hermitian linear systems. In particular, we focus on polynomial preconditioners which are based on information on the eigenvalue distribution of the coefficient matrix. Techniques for obtaining such information adaptively were developed.

Results obtained so far for the case of positive definite Hermitian linear systems indicate faster convergence compared to standard techniques based on the location of the eigenvalues rather than the distribution of the eigenvalues. Moreover, for indefinite Hermitian matrices, we have studied polynomial preconditioners which have the preconditioned linear system indefinite and characterized optimal indefinite polynomial preconditioners. Most of these results are documented in TR89.32. A future technical report on "Pseudo Ritz Values for Indefinite Hermitian Matrices" addresses the problem of estimating eigenvalues of indefinite Hermitian matrices.

Steve Hammond implemented a polynomial preconditioned minimum residual algorithm for a class of complex non-Hermitian matrices on the Connection Machine. The algorithm was tested on 2D and 3D complex Helmholtz equations with the preconditioning approach considerably faster than the non-preconditioning approach especially for 2D problems.

The main focus is now on the development of polynomial preconditioners for general non-Hermitian matrices. As a first step, we have tested techniques for obtaining suitable information on the eigenvalue distribution of non-Hermitian matrices. We plan next to combine polynomial preconditioning with biconjugate gradient type algorithms. The resulting iterative schemes will be tested for nonsymmetric linear systems which arise in CFD. [RIACS TR89.32, TR89.33.]

- **Exact Solution of Linear Equations by Modular Arithmetic on a MIMD vs a SIMD Machine**

Carl Williams

The feasibility of solving a system of linear equations by modular arithmetic on parallel computers was studied. A SIMD approach was implemented on the STARAN and a MIMD approach on the Encore Multimax. A performance analysis was done on the timing of the programs.

## *Parallel Algorithms*

- **Time Dependent Partial Differential Equations**

Yousef Saad, E. Gallopoulos/University of Illinois

Progress has been made on solving parabolic equations by the method of lines combined with Krylov subspace methods. The methods are based on approximating the product of the exponential of a matrix by a vector. Recent work has enabled us to treat time-dependent forcing terms by judicious use of quadrature formulas. Initial results are encouraging. Numerical tests with such schemes have been competitive with explicit schemes as well as the Crank-Nicolson implicit schemes, using Conjugate Gradient method. For problems with constant coefficients, we could show that the method provides very large speed-ups when large dimensional subspaces are taken: a factor of 200 can be gained if a dimension 70 is used as opposed to 6. These methods are essentially explicit in nature in that they do not require solving linear systems. This approach can thus provide a systematic way of generating such explic-

it schemes. It is planned to extend these algorithms to nonlinear systems and explore their use for the solution of the equations arising in CFD. We plan to investigate next the more general time-dependent case where the operator also depends on time, as well as, the nonlinear case (e.g., Burger's equation). [RIACS TR89.19.]

- **Higher Order Solution of the Euler Equations on Unstructured Grids Using Quadratic Reconstruction**

Paul Frederickson, Tim Barth/Code RFTC

We have found an effective definition of a  $k$ -exact reconstruction operator,  $R$ , on an unstructured grid and have applied it in an unstructured-grid flow solver which is, like the structured grid flow solver of Colella and Woodward, higher-order accurate. The reconstruction operator  $R$  is constructed to have two essential properties: The first guarantees that  $R$  exactly conserves the cell averages of flow quantities, while the second requires that when  $R$  is given the cell averages of a degree- $k$  polynomial it will reconstruct that polynomial exactly. Construction of  $R$  requires (a pre-processing step) the solution of a low dimensional least squares problem at every cell.

- **Independent Time Stepping Methods**

Richard Sincovec, Niel Madsen/LLNL

Solving complex physical systems described by time dependent partial differential equations is a compute intensive process even for parallel computers. Often the most interesting part of the solution is a small fraction of the physical domain under consideration. For example, in the vicinity of shocks a small time step is necessary for an accurate and stable solution, whereas, zones far from the shock could be advanced with a larger time step. The Independent Time Step Method (ITSM) was developed as an approach for advancing each zone independently with a time step appropriate to the zone based on accuracy and physical constraints for the problem. On parallel computers, the ITSM is expected to minimize global communication and synchronization. This research examines issues and strategies for implementing the ITSM on parallel computers. Computational results have been obtained that indicate that ITSM is a viable approach for solving time dependent PDEs.

- **Solving the Shallow Water Equations on the Cray and Connection Machine**

Paul Swarztrauber, Richard Sato/NCAR

This work is the most recent in a continuing effort to develop weather forecast methods that are suitable for parallel computers. It is worth emphasizing that one can not simply move an existing code to the Connection Machine but rather one must start from scratch by redesigning basic algorithms so that they are optimal for the parallel architecture. For example, the FFT was completely redesigned for the Connection Machine. This project is an example of the link that exists between NCAR and NAS efforts since both are interested in establishing the merits of parallel computing, both have selected the Connection Machine as an applicable architecture for their computing requirements and both will be using parallel version of algorithms such as the FFT. [RIACS TR89.48.]



- **A Parallel Algorithm for Computing the Eigenvalues of a Symmetric Tridiagonal Matrix**

Paul Swarztrauber

The current state of the art for computing the modes of any dynamical system modeled by a large sparse system of equations begins with the Lanczos method which reduces the large system to a tridiagonal system of equations. The QR algorithm is then used to compute the eigenvalues of the tridiagonal system which also yields the modes of the dynamical system. However, the QR algorithm is not suitable for implementation on a parallel computer. We developed an algorithm specifically for parallel computers that has the lowest complexity (computing time) of any existing algorithm. The computational results demonstrate that the algorithm is also more accurate than the QR algorithm. [RIACS TR89.49.]

- **Ordered Fast Fourier Transforms on a Massively Parallel Hypercube Multiprocessor**

Paul Swarztrauber, Charles Tong

The gap that exists between peak and expected performance of supercomputers continues to widen, primarily due to internal communication requirements that are not included in the popular performance metric of Gflops. Efficient communication can be achieved through algorithmic means like efficient computation. Indeed communication efficiencies obtained by parallel algorithmic means can be several orders of magnitude better than those achieved by routers. These considerations led to the development of a FFT for the Connection Machine that focuses on parallel algorithms for the communication tasks required by the FFT. The resulting algorithm takes advantage of the fact that many different orderings can be implemented with equal ease via the specification of geometries and priorities. Computing time is one third that of previous ordered FFT implementations. [RIACS TR89.50.]

- **Efficient Detection of a Weak Continuous Wave Signal with Linear Frequency Drift**

Paul Swarztrauber, David Bailey/NAS

The purpose of this project was to design an efficient algorithm for detecting signals of the type thought to be the most probable form of initial communication with extraterrestrials. At the heart of the project is the detection algorithm which must be able to identify a weak signal of the type that would be received between two locations that are rapidly accelerating with respect to one another. As a bonus, it appears that this work is also applicable to situations with less acceleration but which require a very high degree of accuracy, such as satellite based aircraft navigational systems and similar systems that are currently being proposed for automobiles. A variant of the FFT was used to speed the computations and provide accurate trigonometric interpolation at the same time. Tests were conducted that demonstrated the ability to detect signals as weak as 1/56th the noise with an average Z-score of 7.2 on a single processor Cray-2 in 2 seconds. [RIACS TR89.51.]

- **Fast Fractional Fourier Transforms and Applications**

Paul Swarztrauber, David Bailey/NAS

Fractional Fourier transforms are defined as transforms that are based on fractional roots of unity rather than the traditional integer roots of unity. Natural cyclic phenomena does not necessarily occur with integer cyclic period but rather periods that have a fractional part with respect to the sample period. Fractional transforms can be used to discover these periods in an efficient manner in a number of problems. [RIACS TR89.52]

- **Particle Simulation Methods on Parallel Computers**

Leo Dagum

On the Connection Machine, we implemented a very efficient direct particle simulation algorithm for two dimensional hypersonic rarefied flows. The implementation is capable of simulating up to 4 million hard sphere diatomic molecules using 64k processors with a performance better than that of a similar, fully vectorized implementation using a single processor of the Cray-2. A performance measure for particle simulations is the average time to advance one particle through one time step. Extrapolated to a 64k machine, the latest performance figure for the Connection Machine is 1.0 microsec/particle/timestep. By comparison, the fully vectorized implementation of this algorithm on a single processor of the NAS Cray-2 can simulate up to 10 million particles but takes 1.8 microsec/particle/timestep. It should be noted, however, that the Cray-2 implementation is 3 dimensional and is more general in scope than the Connection Machine implementation.

An exciting capability of the Connection Machine implementation which cannot be reproduced on the Cray-2 is the real time visualization of a flow. With the frame buffer it is possible to visualize the motion of the particles during a calculation. In addition, using the data vault it is possible to save the images created and play them back at a faster rate. This feature greatly extends the method of particle simulation by allowing for the first time the investigation of transient or unsteady flow phenomena. A video has been created which demonstrates this ability by displaying the simulation of supersonic flow into a resonance tube. [RIACS TR89.44 ]

- **Optimal Moving Grids for Time Dependent Partial Differential Equations**

Andy Wathen

Various adaptive moving grid techniques have been proposed for the numerical solution of time dependent partial differential equations. The precise criterion for grid motion varies, but most techniques attempt to give grids on which the solution of the partial differential equation can be well represented. This project investigated moving grids on which the solutions of the linear heat conduction and viscous Burgers' equation in one space dimension are optimally approximated. Results obtained are for numerical calculations of optimal moving grids for piecewise linear finite element approximation in the least squares norm. [RIACS TR89.42]

- **Earth Mantle Convection on a Highly Parallel Hypercube**

Paul Frederickson, Christopher Lewis, John Baumgardner/LANL

An octahedral triangulation of a two-sphere, with a total of 32,000 triangles, has been developed using a data structure which facilitates fast evaluation of finite element operators. This will allow us to model the convective flow within the mantle, the region which begins a few kilometers below the surface of the earth and extends down about 3000 kilometers. We are currently testing and debugging the operator using surface harmonics, the eigenfunctions of the Laplacian on the two-sphere, and are attempting to develop a version using the fast power-of-two communication routines that are essential for multigrid solution algorithms.

- **Iterative Methods for Markov Chain Modeling**

Bernard Philippe/IRISA, W. J. Stewart/NCSU, Youcef Saad

The purpose of this work was to conduct an in-depth testing and comparison of a number of promising iterative methods for computing the stationary probabilities of Markov Chain. The main linear algebra problem is to compute an eigenvector of a sparse, usually non-symmetric matrix associated with a known eigenvalue. It can be also be cast as a problem of solving a homogeneous singular linear system. These problems tend to be very large and sparse. Another characteristic is that they tend to have clustered eigenvalues that make them difficult to solve. The methods we have tested are based on combinations of Krylov subspace techniques, single vector power iteration and relaxation procedures, and acceleration techniques. The comparison showed that the simpler techniques, such as SOR, are not reliable. Among the most robust techniques we found the preconditioned Arnoldi and GMRES techniques. However, we also found that the more accurate preconditioners tend to be more effective. The work was presented by B. Stewart and partly by Y. Saad at the first Workshop on Numerical Methods for Markov Chains, North-Carolina, Jan 8-12, 1990. [RIACS TR89.39, TR89.40]

- **Computing Radar Cross Sections**

Niel Madsen/LLNL, Alex Woo/RAC

The purpose of this effort is to develop new capabilities for computing radar cross sections. Madsen's code for solving Maxwell's equations in 2-D and 3-D was used. This code is a non-orthogonal finite volume code along with appropriate mesh generation and visualization components. After successfully porting the code to a NAS computer, Alex Woo has continued to use it in his research project.

- **Parallel Sorting Algorithms**

George Adams

A key issue in the design of parallel particle motion algorithms is how to assign the particle data to the different processors. This project investigates the potential of parallel sorting algorithms, particularly algorithms optimized to sort nearly ordered data sets, to restore order to the particles after they have moved one step from their previous, sorted positions. Dynamic particle assignment algorithms could be

based upon such sorting algorithms thereby improving performance by preventing the need for a tradeoff between particle to processor locality and spatial region to processor locality.

- **Efficient Solution Techniques for Problems Using Finite Element and Finite Volume Discretization Algorithms on Hypercube Architectures**

Richard Sincovec, Neil Madsen/LLNL

The objective is to determine efficient ways to solve time-dependent partial differential equation problems on the hypercube machine with non-orthogonal, unstructured grids using either a finite element or finite volume discretization.

The new effort is to determine efficient ways to use the hypercube machine to solve time-dependent partial differential problems which use either finite element or finite volume methods to form the approximating discrete equations. We will assume that the underlying grids are non-orthogonal and unstructured. Initially, our interest will be directed at solving transient problems in a time-accurate manner.

One of the important issues to be addressed is that of developing a strategy to distribute the work to the various processors in a manner that balances the computing load and minimizes the inter-processor message traffic. We expect that some form of domain decomposition method will be appropriate. Since the underlying grids will be unstructured, we will have to develop general data structures that are flexible, efficient and compatible with the MIMD message passing environment. It is felt that this will be the most difficult part of the problem. After addressing the initial time-synchronized problem, we intend to consider the possible application of asynchronous independent time step ideas to further reduce the algorithm communication and synchronization.

Initially, we will consider the solution of Maxwell's equations in two and three dimensions. These equations describe the propagation and interaction of electromagnetic waves with scattering objects. They are used to determine electromagnetic quantities such as the radar cross section of an object. We will consider the use of a finite volume discretization method developed recently by Madsen and Ziolkowski. We will begin our investigations in three dimensions on a logically regular grid and then proceed to unstructured grids in two and three dimensions.

## *Architectures*

- **Sparse Matrix Benchmark**

Youssef Saad, Harry Wijshoff/Illinois

The objective of this benchmark code is to give a realistic measure of how modern supercomputers perform on sparse matrix computations. The emphasis is not only on measuring performance, as is usually done, but also on understanding the difficulties that the computer may be encountering with these very important types of computations. The code is about 10,000 lines long and has been written in ANSI Fortran. It was scheduled for release in first quarter 1990. A presentation was made at the

SIAM conference on parallel processing in Chicago in December, 1989, that showed the first results of the benchmark, including runs on a number of supercomputers. For example looking at the performance loss due to indirect addressing, it was found that the CRAY XMP and CRAY YMP are very much alike, while the CRAY-2 has a distinctly different behavior. Another rather surprising finding was that the effect of memory speed is far more important for (irregular) sparse computations than for regular computations. Thus, on some kernels, the CRAY-2 with static memory (55 ns) could be 1.5 to 2 times as fast as the CRAY-2 with dynamic memory (80ns). Often sparse computations tend to be memory bound rather than CPU bound.

- **Teraflops Computer**

Robert Schreiber, Richard Sincovec

NASA seeks a teraflops computer by the mid 1990s to enable the simulation of full-body transonic flow and the coupling of different physical models. A group of researchers at the IBM T. J. Watson Research Center are developing an ultra-high performance, massively parallel scientific computer, called the Vulcan (formerly the TF-1). During 1989, several meetings were held involving NASA, IBM and RIACS scientists. The purpose of these meetings was to exchange information and to develop the foundations for further collaboration especially with respect to the operating system, language, and applications in computational fluid dynamics.

- **Assessment of Supercomputers**

Robert Schreiber, Horst Simon/NAS and CSC

This is an on-going effort to participate with NASA in assessing the current and projected state of Cray-class supercomputers and parallel research computers including software. An assessment of the Connection Machine was undertaken in 1989 for release first quarter of 1990.

## *Usability and Environments*

- **The Paradigm Compiler**

Jack Dennis, Eric Barszcz/NAS

The Paradigm Compiler is an experimental compiler that will support the programming of numerical scientific codes for the Connection machine. A source program for Paradigm is a collection of function modules written in the SISAL language. The compiler consists of four components: a Parser, a program Analyzer/Transformer, a Code Constructor, and a User Interface. The compiler performs a global analysis of a given linked set of function modules, and synthesizes machine code using advice provided by the user. For the initial version of Paradigm, the goal has been to implement enough of the planned Paradigm Compiler to permit evaluation of its capability for a benchmark Large Eddy Simulation (LES) code provided by NAS. During 1989 work on the compiler reached the point that several simple programs have been compiled and run on the Connection Machine. Preliminary analysis indicates that Paradigm

will achieve performance competitive with other high level languages for the Connection Machine, Connection Machine Fortran in particular, while offering the advantage of its user interface. [RIACS TR89.15]

- **Solving Unstructured Problems on Massively Parallel Computers**

Steve Hammond, Robert Schreiber

The objective is to significantly reduce the communication time in implicit and explicit solution techniques for problems arising from irregular discretizations being solved on massively parallel computers. We assign vertices of unstructured grids to nodes of a hypercube so that neighboring vertices are close to one another in the hypercube.

Currently, the most time consuming part of solving unstructured discretizations of PDE's in 2 and 3 dimensions on massively parallel computers is the communication time. On the Connection Machine, for example, if the communication pattern of an algorithm does not exactly match the interconnection topology of the machine (a hypercube) then one must use the router to exchange data between processors. This is very slow since the router is between 1/6 and 1/30 the speed of the nearest-neighbor network.

We have developed heuristics for assigning the nodes of irregular discretizations of 2D PDE's to the processors on the Connection Machine. Communicating processes are assigned to nearby processors so that we can use the "communication compiler" (newly developed by TMC) and avoid using the router. The communication compiler alone reduces the communication time by a factor of 10 by using many wires in parallel. This will greatly accelerate the rate for solving unstructured problems. We first form an initial guess by mapping a 2D discretization to a regular grid and then use a binary gray-code of the dimensions to assign vertices of the discretization to nodes of the hypercube. Next, we use pairwise exchanges within the hypercube to improve the initial guess. In the cases tested, an unstructured triangular grid around a 3-element airfoil, we reduced the communication time by an additional factor of 3 compared to using the communication compiler alone, for a 30-fold reduction on the CM-2. Although we have focused on mapping to hypercubes there are natural extensions to grids and other interconnection topologies. We are collaborating with Tim Barth (NASA code RFTC) to develop an efficient unstructured CFD code for the Connection Machine based on our graph mapping results. [RIACS TR90.22.]

- **A Massively Parallel Euler Solver for Unstructured Grids**

Steve Hammond, Tim Barth (NASA Ames, code RFTC)

Under the DARPA CAA project, Hammond and Tim Barth have developed a flow code based on work of Barth and Jespersen for solving the Euler equations on triangular unstructured meshes that is very well suited to massively parallel implementation and currently runs on the Connection Machine. A mesh-vertex upwind finite-volume flow solver is used. In traditional parallel and vector mesh-vertex implementations the conserved variables are stored at the vertices of the mesh and computation of flux is done at the edges of the control volumes surrounding each vertex. We call this method of partitioning the problem (*processor-per-edge*). The partitioning

we use assigns a processor to each vertex of the mesh and all computations and data storage are done at the vertices. Computation of the flux function is done at the vertices rather than by edge. We call this method of partitioning the problem (*processor-per-vertex*). There are advantages to using a processor-per-vertex rather than a processor-per-edge approach on a data parallel machine -- the communication is reduced by 50%. (In the processor-per-edge scheme every vertex must send and receive data for every incident edge whereas in the processor-per-vertex scheme, each vertex sends data for every directed edge pointed toward it. At most one vertex communicates for every edge of the mesh which is half as many communications as in the processor-per-edge scheme.)

In addition to reducing the communication in half, we are able to evenly distribute the amount of computation done by each processor. This is very important on a SIMD computer. Two vertices of the mesh share an edge of the control volume. The edges of the mesh are directed to determine which vertex in the pair computes the flux for the common edge. When there is a directed edge from some vertex  $i$  to some vertex  $j$ , then the processor at vertex  $j$  sends its conserved values to the processor at vertex  $i$ , and the flux is computed. A theorem of Chrobak and Eppstein states that for a planar graph (2D mesh) there exists an orientation of the edges such that each vertex has at most three edges directed out from it. Their proof is an algorithm that constructs the orientation in linear time. We are using such an orientation of the edges to guarantee optimal load balancing.

We will continue to develop a general purpose code for the Euler equations using unstructured grid discretizations, in both two and three dimensions. The code will run on the CM and later on the Intel.

- **Machine Independent Programming Languages for Highly Parallel Computers**

Walter Tichy

Parallel programs should be independent of memory organization and communication network, the number of physical processors, and the control mode of the parallel computer. This project defined extensions to Modula-2 for writing highly parallel, portable programs meeting these requirements. The extensions are synchronous and asynchronous forms of a "forall" statement and statements that control the allocation of data to processors. The extensions are small enough to be easily integrated into other imperative languages such as Ada. [RIACS TR89.34.]

- **Sparse Distributed Memory Simulator**

David Rogers

A simulator for Kanerva's sparse distributed memory for the Connection Machine has been completed. This effort and related efforts are described in detail in the Learning Systems portion of this Annual Report.

The research work in computational chemistry is being performed in collaboration with Dr. H. Partridge and Dr. J. Stallcop of NASA code RTC and is a long-term effort to determine the non-equilibrium, high temperature transport properties of air.

Extensive computation were made on nitrogen and oxygen ion-atom interactions including consideration of charge exchange. The work on charge exchange cross sections is entirely new and represents an important scientific result. All the calculations have been completed and the results compare very well in those few regions where there are experimental results. Our findings also exhibit the characteristic fine-structure oscillations predicted by quantum mechanics and hence are of considerable theoretical importance. A paper is in preparation to be submitted to the Journal of Chemical Physics.

Notification was received of acceptance of a paper on chemical education (jointly with Dr. James Eberhard, University of Colorado) entitled, "Simplified Half-Life Methods for the Analysis of Kinetic Data" to appear in the Journal, Education in Chemistry.

- **Super-Workstations and Theoretical Chemistry**

Eugene Levin

During calendar 1989, Levin's technical research activities were primarily concentrated in two general areas: super-workstations and theoretical chemistry. The studies on advanced workstations resulted in a presentation at a NATO workshop on supercomputing in Trondheim, Norway in June of 1989. This presentation was later prepared as a RIACS technical report and will appear as a chapter in a book to be published by Springer-Verlag in 1990. In November of 1989, Levin was invited by NCAR to participate in a workshop to determine the feasibility of using a distributed network of powerful workstations for supercomputing calculations.

The research work in computational chemistry is being performed in collaboration with H. Partridge and J. Stallcop of NASA code RTC and is a long-term effort to determine the non-equilibrium, high-temperature transport properties of air. During 1989, a paper on "Collision Integrals and High Temperature Transport Properties for N-N, O-O, and N-O" was revised and subsequently accepted for publication by the Journal of Thermophysics and Heat Transfer. Extensive computations were made on nitrogen and oxygen ion-atom interactions for a planned future paper dealing with the transport properties of these constituents of high temperature air.

### *Technology Transfer*

The Parallel Systems has attained a high level of visibility and respect in the research community via numerous technical reports, published papers, and professional presentation.

The Parallel Systems has established collaborative efforts with numerous organizations including IBM, Lawrence Livermore National Laboratory, NCAR, Purdue University, University of Illinois, North Carolina State University, Stanford University, Sandia National Laboratory, Los Alamos National Laboratory, University of Colorado, UCLA and Thinking Machines, to name a few.



Parallel Systems has been an active participant as well as host for the Bay Area Connection Machine User's Group. These meetings provide opportunities for sharing of experiences using the Connection Machine.

The Parallel Systems creates, organizes and hosts a weekly seminar series in cooperation with the NAS Systems Division. During 1989 and the first half of 1990, numerous seminars were presented. The success of this seminar series has inspired the creation of a distinguished speakers colloquium series for 1990 on "Parallel Computing in the 90's." To date, Stephen Lundstrom, Burton Smith, Ken Kennedy and Justin Rattner have made presentations in this series.

The Parallel Systems of RIACS was a co-sponsor with Intel Scientific Computers and the Office of Naval Research of the conference Parallel CFD: Implementations and Results Using MIMD Computers which was held in Los Angeles in May, 1989. The Parallel Systems of RIACS was also a co-sponsor with NASA, INTEL, IBM and IU-PUI of the conference Parallel CFD '90: Implementations and Results which was held in Indianapolis in May, 1990.

Roland Freund was awarded the Heinz-Meier-Leibniz award in applied mathematics by the German Secretary of Education for his contributions to applied mathematics.

Roland Freund organized a mini symposium on "Iterative Methods for Solving Linear Systems on Parallel Machines" at the 1989 SIAM Annual Meeting, July 1989, San Diego.

Yousef Saad is completing a technical report that will serve as preliminary documentation for SPARSKIT, a basic tool kit for sparse matrix computations. The objective of this effort is to provide a package for manipulating and working with large sparse matrices. We are about to release the initial version of SPARSKIT. A very limited release was made in 1989. Initial comments from people such as Phuong Vu of CRAY Research and Billy Stewart of North Carolina State University have been extremely encouraging.

Rob Schreiber and Steve Hammond co-organized the 7th Parallel Circus with Gene Golub at Stanford University on March 30 and 31, 1990. The two-day conference attracted 75 scientists from the US, Canada and Europe interested in parallel numerical computing.

Paul Frederickson, Youcef Saad, Robert Schreiber, Richard Sincovec and Horst Simon (NAS/CSC) presented an all day tutorial on "High Performance Matrix Algorithms for Large Scale Scientific and Engineering Computations" at Supercomputing '89, November, Reno, Nevada. This tutorial was also presented at the NASA Ames Research Center in October and video taped for future viewing.

Yousef Saad is in the process of completing a book on "Large Matrix Eigenvalue Problems: Theory and Algorithms."



# **Biographies**



---

## Biographies

### **PETER J. DENNING, Director (1983)**

Denning came to RIACS in June, 1983, from Purdue University, where he was Head of the Computer Sciences Department since July 1979, a Professor of Computer Sciences since 1975, and an Associate Professor of Computer Sciences since 1972. He spent four years as Assistant Professor of Electrical Engineering at Princeton University after receiving his Ph.D. from MIT in 1968. He received an MS degree from MIT in 1965 and a BEE degree from Manhattan College in 1964. His research interests have included operating systems, computer systems architecture, parallel systems, networks, and analytic performance modeling. He has written over 170 technical articles, 100 technical reports, and two books in these areas since 1967. He supervised 13 Ph.D. theses at Purdue. He served as President of the Association for Computing Machinery (ACM) 1980-82 and Vice President 1978-80. He was Editor-in-Chief of the ACM Computing Surveys 1977-78; he is the consulting editor for the MIT Press Series in Computer Science; and he has been Editor-in-Chief of the ACM Communications since 1983. He writes "The Science of Computing" column for American Scientist in each of its six annual issues. He is a Fellow of the IEEE and of the American Association for the Advancement of Science (AAAS). He holds two best-paper awards. He is a member of the Technical Advisory Board for Sequent Computer Corporation and the Research Board for Hewlett-Packard Corporation. He received an honorary Doctor of Laws degree from Concordia University in 1984 and an honorary Doctor of Science degree from Manhattan College in 1985. He received the Computing Research Association's award for service to computing research in 1989, and the Association for Computing Machinery's Distinguished Service Award in 1990.

### **FRANCES B. ABEL, Executive Assistant (1988)**

Abel joined RIACS in July of 1990, transferring from Universities Space Research Association (USRA) at the NASA Marshall Space Flight Center (MSFC), Huntsville, Alabama. She joined USRA in May of 1988, working as Administrative Assistant and Technical Editor to the Director of the Atmospheric Sciences Program. In this position she provided technical support for 22 visiting scientists, and administrative support for over 40 visiting scientists on the Atmospheric Sciences, Microgravity Science, and Space Astronomy/Astrophysics Programs. During this time Abel researched and wrote the draft of the USRA Employee Handbook. Positions prior to USRA include: Executive Assistant to the Vice President of Personnel of the Borg-Warner Corporation; Secretary to the Vice President of Tax Administration of Continental Illinois National Bank, (both in Chicago); and Illustrator/Secretary to Scheduling & Planning Dept. of Lockheed Aircraft International, A.G., Jeddah, Saudi Arabia. Abel is continuing her pursuit of a degree.

**GEORGE B. ADAMS III, Visiting Scientist (1989-90)**

Adams joined RIACS in August, 1983 and worked full-time at RIACS until August, 1987. Since then he has been a Visiting Scientist with RIACS and an Assistant Professor of Electrical Engineering with Purdue University. Adams received the BSEE degree from Virginia Polytechnic Institute and State University in 1978, the MSEE degree from Purdue in 1980, and the Ph.D. in electrical engineering from Purdue University in 1984. His research interests include computer architecture, parallel processing, interconnection network design, and parallel processing algorithms. He has published over 20 papers and reports in these areas and holds a patent for the Extra Stage Cube interconnection network. During 1986 he was a lecturer with the Department of Electrical Engineering, Stanford University, and taught graduate courses on computer design. He is a member of Eta Kappa Nu, Tau Beta Pi, Phi Kappa Phi, and Sigma Xi.

**BILL BEHRMAN, Student (1989)**

Behrman received his B.A. from UC Berkeley with majors in Math and Computer Science. He is presently working towards his M.A. in Applied Math from Berkeley and his Ph.D. in Scientific Computation from Stanford University. His research interest are computational fluid dynamics and parallel computation.

**ROBERT L. BROWN, Scientist (1983)**

Brown came to RIACS in June, 1983, from Purdue University, where he was a Research Assistant in the Department of Computer Sciences. He also served as a Graduate Instructor in charge of a computer science course with an enrollment of over 500. He received his Bachelor's degree in Mathematics from the Ohio Wesleyan University in 1975 and his Ph.D. in Computer Science from Purdue University in 1988. His research areas are operating systems, distributed systems, and programming systems. He has written papers on distributed systems and extensions to the virtual UNIX kernel. His current research involves the specification and design of a graphical distributed programming system for a heterogeneous distributed system. Also, he headed the project to develop and distribute the RIACS Concurrent C compiler for the Sequent multiprocessor.

**WRAY BUNTINE, Visiting Scientist (1990)**

Wray Buntine has a Bachelors in Science (majoring in pure and applied mathematics and statistics) awarded in 1979 and a Diploma in Computer Science awarded in 1982 from the University of Queensland. He then worked in industry a few years before taking up a Ph.D. in 1985 at the University of Technology, Sydney under the supervision of Professor Ross Quinlan. The thesis, "A Theory of Learning Classification Rules," was submitted in early 1990. He has consulted for British Aerospace and BHP, and has been a visiting scientist at the Turing Institute, Glasgow, on two occasions. His broad area of interest is mathematical and probabilistic modeling of problems in intelligent systems, and he has published papers in logic programming, expert systems, plausible reasoning and machine learning.

**TONY F. CHAN, Consultant (1986)**

Chan came to RIACS in January of 1986, on sabbatical from Yale University where he was Associate Professor of Computer Science. He received his Ph.D. from Stanford University in 1978, and his M.S. and B.S. degrees from California Institute of Technology in 1973. His research interests include efficient algorithms in large scale scientific computing, parallel algorithms, and computational fluid dynamics. After completing his sabbatical visit to RIACS, he took a position as Professor of Mathematics at UCLA. He continues his association with RIACS by visiting two days monthly.

**PETER CHEESEMAM, Scientist (1985)**

Cheeseman came to RIACS in April, 1985, from SRI International, where he was a Senior Computer Scientist in both the Robotics and Artificial Intelligence Laboratories. Prior to that he was a Professor at the New South Wales Institute of Technology, Sydney, Australia. He received his Ph.D. in 1979 from Monash University in the area of Artificial Intelligence. He received his M.Phil. in Applied Mathematics in 1973 from Waikato University (New Zealand). His research interests include Artificial Intelligence and Automatic control, induction of models under uncertainty, Bayesian inference, expert systems and robotics. He has published a number of papers in these areas and organized workshops on these subjects.

**BARBARA CURLETTE, Resource Analyst (1987)**

Curlette joined the RIACS administrative staff in November, 1987 after working with Sierra Scientific as a Marketing/Sales and Human Resource Administrator. She attended San Jose State University and received a B.S. in Organizational Behavior/Human Resources from the University of San Francisco in 1986. She is currently working on her Masters Degree in Human Resources and Organizational Development at the University of San Francisco. She is responsible for the administration of financial, contractual, and purchasing functions, and acts as the Institute's Personnel Officer.

**DAVID A. CURRY, Research Associate (1988-89)**

Curry came to RIACS in September, 1988 from the Engineering Computer Network at Purdue University, where he was a systems programmer working on Berkeley UNIX kernel enhancements, systems software development, and systems support. He is finishing course work to complete his B.S. degree in Computer Science from Purdue University. Curry's book, *Using C on the UNIX Systems*, will be published by O'Reilly & Associates in early 1989.

**LEO DAGUM, Student (1987-89)**

Dagum received his B.S. in Engineering Physics from Queen's University, Canada in 1985 and his M.S. in Aeronautics and Astronautics from Stanford University in 1987. His research has concentrated in simulation of fluid flows by particle methods. He is presently working on his Ph.D. developing a rarefied hypersonic flow algorithm for the Connection Machine.

**DOUGLAS G. DANFORTH, Scientist (1988)**

Danforth joined RIACS in January 1988 to work with the Sparse Distributed Memory (SDM) project after being at Wang Labs for seven years as section manager of Signal Processing in their Voice Engineering Department. He received his Ph.D. in Education from Stanford University in 1978. His thesis, entitled "Creative Sequential Classification: an Adaptive Approach to Machine Understanding of Continuous Speech," was on the construction of an automatic speech recognition system that could discover regularities in speech waveforms. He earned a M.S. in Statistics in 1974 from Stanford where, for two years, he was a statistics consultant at The Center for Advanced Study in the Behavioral Sciences. Early work in physics began at Stanford where Danforth earned a B.S. in 1966 and continued at the University of Maryland with a M.S. in 1971. At RIACS, Danforth is working to remove many of the artificial constraints found in the field of speech recognition by applying forms of SDM to front-end encoding, phonetic transcription and grammatical dependencies.

**DAVID DAY, Student (1989)**

Day received his B.A. in Mathematics from the UC Berkeley in 1986. He is presently working on his Ph.D. developing efficient preconditioners for elliptic problems and parallel algorithms computing them.

**MARTHA DEL ALTO, Secretary (1990)**

Del Alto joined RIACS in February of 1990. She came to RIACS from Sumitomo Bank, where she held the position of Bank Services Representative. She received her AA degree from Mission College in 1985 and currently working on her BA at SJ-SU on a part-time basis.

**JACK B. DENNIS, Visiting Scientist (1988-89)**

Dennis came to RIACS in May of 1988, as a Visiting Scientist following his career as Professor of Computer Science and Engineering at MIT. He is a graduate of MIT, having received his doctorate in 1958. In his years at MIT, Dennis led the Computation Structures Research Group of the Laboratory for Computer Science for twenty years, supervised more than twenty-five doctoral thesis projects, and developed six subfields of instruction in new areas of computer theory and computer systems. Dennis is known throughout the world for his contributions in the field of computer system architecture. In the 1960's he had major influence on the virtual memory architecture of the Multics system implemented by Project MAC at MIT, work for which he was elected Fellow of the IEEE. The dataflow model of computation was formulated by Dennis and his students in the MIT Computation Structures Group and led to a series of proposals for parallel computers having dataflow architecture. Dennis received the Eckert-Mauchly Award in 1984 in recognition of this work. His current projects concern compiling methods for functional programming languages and highly parallel computers, and the evaluation of dataflow multiprocessor computers for a variety of applications.



**DEE DOYLE, Research Associate (1988)**

Doyle joined RIACS in October 1988. She received a B.S. in Political Science from Arizona State University (1977) and M.S. in Computer Science from New Mexico State University (1987). From 1980-1984 she negotiated and administered grant and contract programs awarded to Purdue University. From 1987-1988 she worked on internal research and development projects at TRW.

**LORRAINE FISHER, Secretary (1988-90)**

Fisher joined the RIACS staff in January, 1988. Prior to that date she was employed as an administrative assistant to the Regional Transportation Manager of the AMPEX Corporation. She is currently in the process of finishing her coursework toward a B.S. degree in Mathematics and Computer Science at San Jose State University. She was the secretary for the Networked Systems.

**PAUL O. FREDERICKSON, Senior Scientist (1988)**

Frederickson came to RIACS in May 1988, from the Computing and Communications Division of Los Alamos National Laboratory. He has been involved in the development of parallel algorithms at Los Alamos since 1979, with emphasis on multigrid algorithms and their application to fluid dynamics and particle models. During the academic year 1985-86 he was a visiting scientist at Christian Michelsen Institute in Bergen, Norway, where he demonstrated the value of highly parallel computers to the offshore oil industry. During the academic year 1978-79 he worked in global atmospheric modeling at Goddard Spaceflight Center, while on leave from Lakehead University in Canada where he had taught mathematics and computer science since 1969. He was a visiting member of the mathematics institute of the Technische Universitaet, Muenchen, Germany during 1973-1974, doing research in the fast solution of very large linear systems. He was an assistant professor of mathematics at Case Institute of Technology from 1964 until 1969, after earning his Ph.D. at Nebraska University in 1964. He is a member of ACM, SIAM and AMS.

**ROLAND FREUND, Scientist (1988-89)**

Freund joined RIACS in May, 1988. He is on leave from the University of Wuerzburg, Federal Republic of Germany, where he holds a position corresponding to that of an Assistant Professor in the United States. He received his Diploma in Mathematics in 1982 and his Ph.D. in Mathematics in 1983, from the University of Wuerzburg. From September 1985 to September 1986 he was working as a Visiting Research Associate in the Computer Science Department at Stanford University. His research interests are mainly in numerical linear algebra and in approximation theory and its applications, especially in iterative matrix computations. Recent interests include the solution of linear systems arising in fluid computations and the development of preconditioners for massively parallel architectures.

**MARIA L. GALLAGHER, Networked Systems Projects Coordinator (1986)**

Gallagher joined the RIACS staff in April, 1986. Prior to that time she had been employed in a variety of administrative positions, as a teacher, and as a computer consultant. She received her A.A. in Political Science from the College of San Mateo, and her BA in History and Political Science, as well as a Lifetime Teaching Credential, from San Jose State University. Graduate work on her combined MA (Art Education/Child Psychology) and Early Childhood Teaching Credential has been completed at San Jose State University. She has completed additional course work in Math, Engineering, Networks and Computer Science towards a degree in Computer Science. During 1988, Gallagher acted as Projects Coordinator for the Telescience Testbed Pilot Program and the Telescience Applications Project Areas. She continues to provide the administrative support for Ames personnel involved in the Telescience area. In addition, Gallagher is working with the NASA Science Internet Project Office at Ames to assist them in providing user support services and documentation and in defining and maintaining user communications requirements.

**DEANNA M. GEARHART, Administrative Secretary (1988)**

Gearhart joined the RIACS staff in May, 1988. Prior to that time she worked as accounts payable manager, purchasing agent and contract administrator for Martin J. O' Sullivan Construction Inc.

**DAVE GEHRT, Senior Research Associate (1988)**

Gehrt rejoined RIACS in January, 1988, as a Research Associate, after a two and one-half year absence during which he was employed by the Antares Group, Inc. of San Diego. He served as RIACS Facilities Manager from April 1988 through October 1989. Prior to his employment at Antares, he was the Facilities Manager at RIACS from January, 1984, until July, 1985. Earlier he had been employed for seven years by the SYSCON Corporation in a variety of positions ranging from programmer to Operations Manager. Prior to his employment by SYSCON, he practiced law in the State of Washington, taught Jurisprudence at Shoreline Community College in Seattle, and was employed as a programmer/analyst by Mellonics, a division of Litton Industries, and the Commercial Airplane Division of the Boeing Company. He received a B.A. from Washington State University in Pullman, Washington, in 1965, and a J.D. from the University of Washington in 1972.

**JULIAN E. GOMEZ, Consultant (1988-89)**

Gomez came to RIACS as a Scientist in January, 1986 from Cranston/Csuri Productions Inc., where he was director of R&D. In January 1988 he left to pursue his entrepreneurial interests in animation software and remained a consultant for the RIACS scientific computing environments project. He received his Ph.D. in Computer and Information Science from Ohio State University in 1985. Prior to his graduate program he spent four years in the Computer Graphics Laboratory at the Jet Propulsion the Jet Propulsion Laboratory. He received an A.B. with honors from the University of California at Berkeley in 1977. His areas of interest are 3-

D computer animation, chaos and natural phenomena and he has published several technical papers in these areas. He is a member of ACM, SIGGRAPH, and the IEEE Computer Society. He is Chairman of the Bay Area ACM/SIGGRAPH Technical Interest Group on Performance Evaluation. He is a member of the ARC Code R Computer Graphics Planning and Review Committee.

**RICHARD F. HAINES, Scientist (1988)**

Haines joined RIACS in March, 1988. Before this he had recently retired from NASA Ames as Chief of the Space Human Factors Office where he directed numerous projects for the Space Station Program. Before then he was a research scientist in various divisions at Ames working on aeronautical (head up display, aircraft cockpit displays and simulations) and astronautical projects (Gemini, Apollo, Shuttle). He received the M.A. (1962) and Ph.D. (1964) from Michigan State University in experimental psychology. He has published over 65 journals articles, over 35 government technical reports, three books, and holds four US Patents (one pending). He has also produced two technical movies for NASA. He served as a contributing editor to *Aviation, Space and Environmental Medicine*, and *Kronos-Journal of Interdisciplinary Synthesis*. He is currently an Associate Fellow of the Aerospace Medical Association and founding chairman of its Space Station Human Factors subcommittee; a full member of the Society of Automotive Engineers (and member of its S-7 aviation subcommittee); a member of the International Society of Air Safety Investigators (and director of the San Francisco Chapter); a member at large of Sigma Xi; and a member of the Optical Society of America. His full biography may be found in *Who's Who in America* and *Who's Who in the West*, *American Men and Women of Science*, and other references.

**STEVEN W. HAMMOND, Research Associate (1988)**

Hammond joined RIACS in October, 1988, from Rensselaer Polytechnic Institute in Troy, New York. He is currently enrolled in the Ph.D. program in Computer Science there. He received his B.A. in Mathematics in 1983 and his M.S. in Computer Science in 1984, from the University of Rochester in Rochester, New York. From 1984 until 1987 he was on the research staff at the Corporate Research and Development Center of General Electric Co. in Schenectady, New York. While with GE he conducted research in the area of algorithms and special purpose architectures for computational numerical linear algebra. He filed seven patent applications concerning a method and apparatus for solving large, sparse linear systems on linear arrays of processors. In 1987 Steve left GE to return to school to obtain his Ph.D. at Rensselaer. His research is in the area of computational numerical linear algebra for large scale scientific calculations. He has been awarded an IBM Doctoral Fellowship for the 1988-1989 academic year. He is a member of SIAM and Sigma Xi.

**KATHRYN HAWKEN-CRAMER, Executive Secretary (1987-89)**

Hawken-Cramer joined the RIACS staff in June, 1987. Prior to that she worked as sales office manager for Robinson-Nugent, an electronic components firm and as hostess/manager at Scott's Seafood Restaurant in Palo Alto. She studied chemistry at the University of Arizona from 1974 to 1977.

**MICHAEL HOWELL, Research Associate (1988-90)**

Howell joined RIACS in August, 1988 from the Tactical & Training Systems Division of Logicon Inc. At Logicon he was a staff hardware engineer on testbed and simulation projects. Before Logicon, he worked as an engineer performing testing, modification and training on Naval shipboard combat weapons systems. He spent over a decade in the US Navy. During that time, he was either in training for, or performing in, roles as electronic technician, instructor, and nuclear reactor operator. His principle current interests are systems integration and networks.

**LOUIS JAECKEL, Scientist (1987)**

Jaekel came to RIACS in August 1987 to work with the Sparse Distributed Memory group. He received his B.A. in Mathematics from UCLA in 1961, where he won the prestigious William Lowell Putnam Prize Fellowship to Harvard University - a mathematics competition open to all undergraduates throughout the United States and Canada. He received his M.A. in Mathematics from Harvard University in 1962 and received his Ph.D. in Statistics from U.C. Berkeley in 1969. Also in 1969 Jaekel was awarded the Bernard Friedman Memorial Prize for best dissertation in a mathematical subject at U.C. Berkeley. He was a member of the technical staff at Bell Laboratories from 1970 to 1972, and an assistant professor of statistics at U.C. Berkeley from 1972 to 1977. Jaekel has a Harvard law degree. From 1977 to 1987 he taught statistics and computer science at U.C. Berkeley and at Mendocino College and consulted in the areas of legal research and forensic statistics. His interests include mathematics, statistics, and computer science.

**MARJORY J. JOHNSON, Senior Scientist (1984)**

Johnson came to RIACS as a Scientist in January, 1984 from the University of Missouri at St. Louis where she was an Associate Professor of Computer Science. She was promoted to Senior Scientist in 1988. She received a BA Degree in Mathematics from Florence State University and M.S. (1968) and Ph.D. (1970) degrees in Mathematics from the University of Iowa. From 1970-1976 she was an Assistant Professor of Mathematics at the University of South Carolina. From 1976-80 she was a systems analyst at NCR Corporation, where she worked with microcomputer systems and computer communication networks. She has published technical papers in both mathematics and computer science.

**RICHARD G. JOHNSON, Visiting Scientist (1988)**

Johnson joined RIACS as a visiting scientist in January, 1988. Prior to that he was Acting Science Advisor to the President and Acting Director of the White House Office of Science and Technology Policy (OSTP), Executive Office of the President from May to October 1986. Johnson was Assistant Director for Space Science and Technology in OSTP from November 1983 to October 1987. He received his B.S. in Physics from Antioch College in 1951 and his Ph.D. in Physics from Indiana University in 1956. That same year, he joined the Lockheed Palo Alto Research Laboratory of the Lockheed Missiles and Space Company, where, for 27 years, he conducted a broad range of research in low energy nuclear physics and in the space

sciences. He was Manager of the Space Sciences Laboratory for ten years and Senior Science Advisor to the Director of Research for five years. He was a visiting professor at the University of Bern in 1980.

**CHARLES C. JORGENSEN, Visiting Scientist (1990)**

Jorgensen joined RIACS in 1990 as a Visiting Scientist in the Learning Systems. He received his Ph.D. in Mathematical Psychology in 1973 from the University of Colorado and was a Visiting Fellow at Carnegie Mellon University in 1974. Dr. Jorgensen spent nine years as a team chief of Air Defense Systems at the Army Research Institute, followed by 5 years with Oak Ridge National Laboratories. In 1987 he left to become head of Neural Network Research for Thomson-CSF in Palo Alto, prior to his current position as Chief of the Intelligent Systems Technology branch at NASA Ames. Dr. Jorgensen is the author of over 60 publications and book chapters in the areas of intelligent robotics, neural networks, and human factors engineering.

**ISA JUBRAN, Student (1989)**

Jubran came to RIACS in June 1989 while a M.S. student in Science at Oregon State University. He was a Graduate Teaching Assistant in the Department of Mathematics. His main focus was differential equations, linear algebra, calculus, and college algebra. Jubran earned his B.S. degree majoring in mathematics from Bethlehem University, West Bank, Israel in June 1984.

**PENTTI KANERVA, Scientist (1985)**

Kanerva came to RIACS in October 1985 from Stanford University to establish the Sparse Distributed Memory project. While at Stanford, Kanerva developed a large terminal network, software and hardware for technical writing, typesetting software, algorithms for searching and sorting, new memory architectures, new CPU architectures, and directed the technical development of a campus-wide network for text-handling. He was also a postdoctoral fellow and is presently a visiting scholar. In 1965-1967 he organized the first computer center for the University of Tampere, Finland, and taught mathematics, statistics, and computer programming. Prior to working at Tampere, he worked in experimental design and statistical analysis at the Forest Research Institute of Finland and the Finnish State Computer Center. He has a Ph.D. in Philosophy from Stanford (1984) and M.S. in Forestry from the University of Helsinki (1964). His current research is based on his dissertation on memory and is aimed at building practical computers with some of the desirable properties of human memory. His interests include computer-aided communication by people, the computing power of a network of computers, and the design of systems for people without a strong background in computers.

**KIMO K. KASKI, Consultant (1990)**

Kanski came to RIACS in January 1990 from Tampere University of Technology, Finland, where he was a Professor in Microelectronics. His work centered around Sparse Distributed Memory with Pentti Kanerva. He received his Ph.D. in Theoretical Physics from Oxford University, England, in 1981.

**PHIL KLIMBAL, Consultant (1989)**

Klimbal joined the RIACS staff in 1989 from the Engineering Computer Network at Purdue University where he was a system programmer working on BSD UNIX kernel enhancements and systems support. He received a B.S. in Computer Science from Purdue University in 1985. His areas of interest include operating systems, distributed computing, networking, and system integration.

**ANNE F. KOHUTANYCZ, System Administrator (1985)**

Kohutanycz joined the RIACS staff in January, 1985, after holding the positions of Office Manager and Administrative Assistant at CRS Inc., San Jose, California. An active member of the Air Force Reserves since 1978, she is currently serving as the Command Section Administrator for the 349th Field Maintenance Squadron, Travis AFB, California. She received a B.S. in Business Administration from San Jose State University in 1984 and an A.A. in Computer Science from Hartnell College in 1982.

**BARRY M. LEINER, Senior Staff Scientist (1985)**

Leiner joined RIACS in August 1985, as Senior Staff Scientist, served as Assistant Director for Networked Systems from August 1988 through March 1990, and resumed his position as Senior Staff Scientist in April 1990. He has been formulating and carrying out research programs in distributed systems. These programs range from the development of advanced computer and communications technologies through to the application of such technologies to scientific research. Prior to coming to RIACS, he was Assistant Director for C3 Technology in the Information Processing Techniques Office of DARPA (Defense Advanced Research Project Agency). In that position, he was responsible for a broad range of research programs aimed at developing the technology base for large-scale survivable distributed command, control and communication systems. Prior to that, he was Senior Engineering Specialist with Probe Systems, Assistant Professor of Electrical Engineering at Georgia Tech, and Research Engineer with GTE Sylvania. Dr. Leiner received his BEE from Rensselaer Polytechnic Institute in 1967 and his M.S. and Ph.D. from Stanford University in 1969 and 1973, respectively. He has published research papers in a number of areas, including direction finding systems, spread spectrum communications and detection, data compression theory, image compression, and most recently computer networking and its applications. He received the best paper of the year award in the IEEE Aerospace and Electronic Systems Transactions in 1979 and in the IEEE Communications Magazine in 1984. Leiner is a Senior Member of the IEEE and a member of ACM, Tau Beta Pi and Eta Kappa Nu.

**EUGENE LEVIN, Assistant Director (1983)**

Levin came to RIACS as a Scientist in June, 1983. He served as Deputy Director from August 1987 through March 1990, and then took on the position of Assistant Director in April 1990. Before coming to RIACS, he was Chief Engineer of the Systems and Software Division of System Development Corporation (SDC), where he helped analyze, design, and implement portions of large scale military command, control, and communications systems. Prior to joining SDC, he was Vice President of Culler-Harrison, Inc., where he helped develop an advanced array processor

for signal analysis. During 1961-70 he worked at the Aerospace Corporation as Director of the Guidance and Control Subdivision. Levin has served on the AFIPS Board of Directors and as Chairman of the AIAA Computer Systems Committee. He has over fifty technical publications in various fields of applied mathematics, physics, and computer technology. His education includes an A.B. (1950) and M.A. (1951) in Physics from UCLA, and a Ph.D. (1955) in Mathematics from UCLA. In preparation for his current position, he returned to school and received a B.A. in Chemistry from the University of Colorado in May 1983. He is currently working in collaboration with the Ames Computational Chemistry Branch to determine the high temperature transport properties of air.

**CHRISTOPHER LEWIS, Research Associate (1989)**

Lewis joined RIACS in January, 1989. He has received his B.S. in Computer Science from George Mason University in 1987. From June 1987 to January 1989, he was a consultant at the Naval Research Laboratory's Connection Machine facility, where he worked on the development of applications software for the Connection Machine. Previous to this, he was employed at Comsys Inc., The Pentagon, advising Unisys and Defense Communications Agency personnel on the Ada reimplementation of the NMCC Automated Control Executive (NACE). He is currently working with the Parallel Systems on various projects involving the use of the Connection Machine. This will include the design and implementation of software for the Connection Machine. He will also be developing tools that will assist others in the use of the Connection Machine.

**EGON E. LOEBNER, Associate (1988-90)**

Loebner, an Associate of the SDM project, is Counselor for Science and Technology, Hewlett-Packard Laboratories in Palo Alto. He joined HP in 1961 as manager of Optoelectronics of HP Associates after heading optoelectronics research for six years at RCA Laboratories and prior to that for three years at Sylvania Electric Products, Inc. From 1974 to 1976 he served as Chief of the Science Section at the American Embassy in Moscow. He received a Ph.D. in Physics from SUNY at Buffalo in 1955. His thesis dealt with the electronic properties of carbons and graphite. In 1959 Loebner received an RCA Achievement Award for research on the chemistry of high-band-gap semiconductors. He has authored over 40 patents and 50 papers in bionics, biophysics, chemistry, cognitive science, computational linguistics, electronics, human factors, information displays, material science, optics, physics and telecommunications. Loebner has been elected Fellow of the IEEE in recognition of contributions to optoelectronics materials, devices, and networks, and for transdisciplinary research across life, social, and computer sciences. He is listed in *Who's Who in the West* and in *Who's Who in the World*. His main contributions are the establishment of optoelectronics as a technology, pioneering of LED displays, design of a functional model of the frog retina, and construction of a psychophysical model of discovery and invention. We are very sorry to report that Dr. Loebner died of cancer on December 30, 1989.

**ARMANDO E. LOPEZ, Consultant (1987)**

Lopez joined RIACS in August, 1987, from NASA Ames Research Center where he had worked as a Research Scientist for thirty-five years. His work included nine years in aerodynamic research at the Twelve-Foot Pressure Wind Tunnel, eleven years in Guidance and Navigation and Stability and Control, and fifteen years as a staff assistant to the Flight Systems and Simulation Research Division. He obtained a B.A. in Physics from the University of California, Berkeley in 1951. He has assisted RIACS by simplifying the apparent complexities of the government procurement and financial system.

**DAVID MACKAY, Summer Student (1990)**

David Mackay received his B.S. and M.S. degrees from Brigham Young University in 1986 and 1987. He is working on a Ph.D. at Stanford University in the civil engineering department where he studies computational mechanics. His interests include the solution of structural mechanics problems on parallel computers. He is working on solution of electrodynamic problems in three dimensions on the Intel hypercube.

**NIEL K. MADSEN, Visiting Scientist (1988)**

Madsen began a collaboration with RIACS in October, 1988, with the Parallel Systems. He received his A.B. in Physics and Mathematics in 1964 from the California State University at Chico, and his Ph.D. in Mathematics in 1969 from the University of Maryland. From 1968 to 1972 he worked at the Westinghouse Bettis Atomic Power Laboratory as a numerical analyst and also taught at Carnegie-Mellon University. Since 1972 he has worked at the Lawrence Livermore National Laboratory in the Computations and Engineering Departments where he currently heads an electromagnetics research group. His work has concentrated on problems associated with the numerical solution of partial differential equations in general with specific interests in method of lines techniques, general purpose PDE software, matrix algorithms for vector and parallel computers, adaptive grid solution techniques involving finite difference, finite element and finite volume methods.

**COE MILES-SCHLICHTING, Associate (1988)**

Research Scientist Miles-Schlichting has been a member of RECOM Technologies, Inc. technical staff since April 1987 and, under sponsorship of the NASA Ames Information Science Division, he has worked with the SDM group as associate in March of 1988. He received a B.S. in Electrical Engineering from Oregon State University in 1984 and a M.S. in Computer Science from Santa Clara University in 1987. His work with the SDM Cerebellum group is centered in doctoral research, modeling the mammalian cerebellum. Miles-Schlichting's previous engineering positions include product development engineer in a small laser systems company, and lead designer for the control system of a multiprocessor system designed for signal acquisition and analysis.



**NOEL NACHTIGAL, Summer Student (1990)**

Massachusetts Institute of Technology Nachtigal received his B.S. in Mathematics and Physics (double major) in 1987, from the University of South Carolina. He is currently working on his Ph.D. at the Massachusetts Institute of Technology, under the supervision of Nick Trefethen. His research interest is in algorithms for nonsymmetric linear systems.

**BRUNO A. OLSHAUSEN, Research Associate (1987-89)**

Olshausen joined RIACS in November, 1987, after receiving his M.S. and B.S. degrees in Electrical Engineering from Stanford University. At Stanford he held research assistantships in the Robotics and Computer Vision Laboratory, and in the Perception and Cognition Laboratory (Code FL) at NASA Ames Research Center, where he worked on programs for computer graphics and image processing. During internships at the Hughes Artificial Intelligence Center (1983-1985) and the German Aerospace Research Establishment (1985), Olshausen developed programs for rendering 3-D perspective views of terrestrial and planetary surfaces. Olshausen is currently working on neural-network approaches to visual pattern recognition.

**BARNEY PELL, Student (1989)**

Pell came to RIACS in June 1989 while a student at the University of Cambridge in Cambridge, England and a student associate of the Artificial Intelligence Center of SRI International in Menlo Park, California. He was pursuing his Doctor of Philosophy degree in artificial intelligence. Pell received his B.S. with highest distinction in symbolic systems: Artificial Intelligence from Stanford University, Stanford, California in June 1989.

**JO ANNE M. RAMIREZ, Secretary (1988-90)**

Ramirez joined RIACS as secretary for Learning Systems in December, 1988. She left RIACS in May 1990 to take a position at Sun Microsystems, Inc. She was previously with Hewlett-Packard Headquarters in Palo Alto, where she was Marketing Communications secretary for the Test and Measurement Marketing Group. At HP she also worked for a Group Controller, Market Development Manager, and Sales/Service Manager.

**MICHAEL R. RAUGH, Chief Scientist (1985)**

Raugh joined RIACS as Chief Scientist in January 1985. Raugh founded the Advanced Algorithms and Architectures project (now part of Parallel Systems) and, with Pentti Kanerva, founded the Sparse Distributed Memory (SDM) project. In 1989 he became Assistant Director in charge of the new Learning Systems. Raugh came to RIACS from Hewlett-Packard Laboratories of Palo Alto, where he worked in mathematical applications and chaired organizational efforts leading to establishment of the HP Labs Scientific Computing Facility. He received his M.S. and Ph.D. in Mathematics from Stanford University in 1977 and 1979, respectively. He received a B.S. from UCLA in 1962. A former Fulbright/Hayes and National Science Foundation Fellow in mathematics, Raugh has worked as applications and systems programmer at Lawrence Berkeley Laboratory, US Geological

Survey, and the Institute for Mathematical Studies in the Social Sciences at Stanford University. He has written data management systems, Monte Carlo simulations, computer-assisted-instruction programs, program development systems, and calibration and digital filtering programs. At Hewlett-Packard Raugh invented procedures for calibrating electron-beam lithography systems, based on his theory of self-calibration. His central interest is the application of mathematics and computers to modeling physical and biological systems. Raugh is Project Manager for the Sparse Distributed Memory Project.

**KAREN S. RICKMAN, Secretary (1990)**

Rickman joined the RIACS staff in May 1990. Prior to that date she was employed as an Administrative Assistant for Boeing Aerospace Operations at ARC. In 1984 she received a B.A. in Creative Arts from San Jose State University with a minor degree in English. She is secretary to the Networked Systems.

**DAVID ROGERS, Scientist (1987)**

Rogers joined the RIACS staff in August 1987. He received a B.S. in Chemistry and a B.A. in Applied Mathematics from UC Berkeley, and a Ph.D. from the UC Santa Cruz. Following a one-year postdoctoral appointment at the MIT Artificial Intelligence Lab, he spent three years at the University of Michigan working with Douglas Hofstadter in the Fluid Analogies Research Group. His interests include massive parallelism, artificial intelligence, and artificial neural systems.

**YOUCEF SAAD, Senior Scientist (1988-90)**

Saad joined RIACS in May, 1988. He received the "Doctorat de 3eme Cycle" in 1974, and the "Doctorat d'Etat" in 1983 from the University of Grenoble (France). He was an Associate Professor in the mathematics department and senior computer scientist with the Center for Supercomputing Research and Development of the University of Illinois, at Urbana Champaign, from 1986 to 1988. He held the positions of Research Scientist and then Senior Scientist at Yale University from 1981 to 1983 and from 1984 to 1986. He also held visiting positions at the University of Illinois at Urbana Champaign in 1980 and the University of California at Berkeley in 1981. He is on the editorial board of *SIAM Journal of Numerical Analysis* and is a member of SIAM and the IEEE Computer Society. His research interests include: parallel computing, numerical linear algebra, iterative methods, eigenvalue methods, sparse matrix computations, and control theory.

**CATHY L. SALAZAR, Secretary (1986-89)**

Salazar joined RIACS as secretary for the Parallel Systems in August, 1986. She came to RIACS from Lockheed Missiles and Space Company where she held the position of project secretary in the Space Systems Division. From 1981 to 1986 she was Project Secretary for Bipolar Microprocessor Engineering at Advanced Micro Devices. She received her AA degree from West Valley College in 1981.

**HAYCHAN SARGENT, Executive Assistant (1989-90)**

Sargent's role at RIACS was executive assistant to the Director. Before she came to RIACS, she was an administrative manager at Olivetti Research Center in Menlo Park, California, acting as a liaison between Olivetti US and Olivetti Headquarters in Italy and was responsible for the operations of the Olivetti Software Technology Laboratory. Positions prior to Olivetti include: office manager and personal assistant to President of Acorn Research Center, Inc.; executive assistant to manager of Imaging Laboratory at Xerox Palo Alto Research Center where she also worked closely with authors and publishers to produce and illustrate technical books; member of the Screen Fonts Design Team at Adobe Systems, Inc., 1986-1988 (contractor); research assistant at Stanford Medical Center and Children's Hospital at Stanford. Her artwork can be found in ACM and IEEE journals and the New England Journal of Medicine.

**ROBERT S. SCHREIBER, Senior Scientist (1988)**

Schreiber joined RIACS in August, 1988. He received his A.B. in Mathematics in 1972 from Cornell University and his Ph.D. in Computer Science in 1977 from Yale University. During the years 1977-1979, he worked on computational fluid dynamics at Caltech. This work involved efficient numerical procedures for steady flows and their implementation on the then new vector computers. During the years 1979-1984, he was Assistant Professor of Computer Science at Stanford. While there, he worked on sparse matrix computation, matrix computation, parallel algorithms and systolic arrays, and signal processing. He was also a consultant to ESL, Inc. where he helped to design a parallel computer system for sonar signal processing. In 1984, he left Stanford for the Saxpy Computer Corporation. There he became Vice President and Chief Scientist and was the principal architect of the Saxpy Matrix 1, which was a 1 gigaflop matrix computer. He was involved in the design of successor to the Matrix 1. During 1986-1987, he left Saxpy to take a position as Associate Professor of Computer Science at Rensselaer Polytechnic Institute in Troy, New York. He is currently on leave from this position. He is now supervising two RPI Ph.D. students. His research involves parallel numerical algorithms and parallel computer architectures.

**B. DAVID SEMERARO, Summer Student (1990)**

Semeraro received his B.S. in Aeronautical Engineering from the University of Illinois in 1980 and his M.S. in 1982. He spent three years in the aerodynamics research department at Northrop Aircraft before returning to Illinois to work on a Ph.D. in computer science. His research is concentrated on parallel algorithms for computational fluid dynamics.

**RICHARD F. SINCOVEC, Assistant Director (1988)**

Sincovec joined RIACS in June, 1988. He received his B.S. in Applied Mathematics from the University of Colorado-Boulder in 1964 and his M.S. and Ph.D. in Applied Mathematics from Iowa State University in 1967 and 1968, respectively. Prior to joining RIACS, he was Professor and Chairman of the Computer Science Department at the University of Colorado, Colorado Springs, Manager of the Numerical Analysis Group at Boeing Computer Services, Professor of Computer Sci-

ence and Mathematics at Kansas State University and a Senior Research Mathematician at Exxon Production Research. He has also been affiliated with the Software Engineering Institute of Carnegie-Mellon University, Lawrence Livermore National Laboratory and Hewlett-Packard. His research interests are in scientific computing, parallel computing, numerical analysis, numerical algorithms and the application of modern software engineering and computer science concepts to scientific computing. His most recent work is in adaptive techniques for solving partial differential equations, parallel algorithms and their implementation in Ada, and in computer aided software engineering tools. Prior work includes the development of algorithms and general purpose software for the solution of problems described by partial differential equations and the solution of systems of differential/algebraic equations. He is the coauthor of five books and a member of SIAM and ACM.

**DANNY C. SORENSEN, Visiting Scientist (1990)**

Sorensen came to RIACS from Rice University where he was Professor of Mathematical Sciences. Sorensen received his B.S. in Mathematics from the University of California, Davis, in 1972 and his Ph.D. in Mathematics from the University of California, San Diego, in 1977. He was an Assistant Professor of Mathematics at the University of Kentucky from 1977-1980. He then joined the Mathematics and Computer Science Division of Argonne National Laboratory where he became a Senior Computer Scientist. Sorensen has been a visiting professor at U.C. San Diego and at Stanford University. He was also a visiting scientist for one year at the Center for Supercomputing Research and Development at the University of Illinois, in Urbana. He joined the faculty of Rice University in 1989. Sorensen was one of the founders of the Advanced Computing Research Facility at Argonne. This facility was one of the first to provide public access to a variety of parallel computers and has served as a model for a number of similar facilities that now exist across the nation. His research interests are in computational mathematics with emphasis on numerical linear algebra and parallel computing. He has also worked extensively in the area of nonlinear numerical optimization. Sorensen visited RIACS during July and August, 1990, to implement some parallel algorithms for eigenvalue problems.

**HENRY A. SOWIZRAL, Scientist (1988)**

Sowizral joined RIACS in July of 1988. He received his Ph.D. and M.Phil. in Computer Science from Yale University in 1982 and 1977. He earned his B.S. in Information and Computer Science from the University of California at Irvine in 1975. Sowizral came to RIACS from Schlumberger where he was principal investigator on the Distributed Systems technology project funded by the Office of Naval Research. That project extended and continued the Time Warp Simulation System, a technology he developed at the Rand Corporation for taking a discrete-event simulation and executing it on a network of processors faster than on a single processor. Prior to working with Schlumberger, Sowizral was at the Rand Corporation where he jointly developed the theory for Time Warp and was the principal investigator on a DARPA-sponsored project to develop an English-like expert system programming language called ROSIE.

**ROBERT STEFFEN, Research Associate (1990)**

Steffen joined the RIACS staff in January of 1990. Prior to this he worked five years at NASA Ames Research Center, in space hardware development. He received his B.S. from Rensselaer Polytechnic Institute, Troy, New York in 1986.

**HELEN STEWART, Secretary (1987)**

Stewart joined RIACS in September, 1987. She began as a general secretary, became the secretary for the Sparse Distributed Memory project, and then the secretary for the Learning Systems. Prior to her arrival at RIACS, she was regional secretary for the McDonnell Douglas Corporation. She received her B.S. in Business at Evergreen College.

**PAUL SWARZTRAUBER, Visiting Scientist, (1989)**

Swarztrauber visited RIACS from April 1 through December 31, in 1989, as part of the Parallel Systems. He received his B.S. in Engineering Physics from the University of Illinois in 1959 and his M.S. in Applied Mathematics and Ph.D. in Mathematics from the University of Colorado in 1966 and 1970, respectively. His permanent position is Senior Scientist at the National Center for Atmospheric Research in Boulder, Colorado. His research interests include FFTs and parallel algorithms.

**EVANGELINE TANNER, Secretary (1990)**

Tanner joined RIACS as Secretary to the Parallel Systems in April 1990. A few years prior to her joining RIACS, she held such positions as Secretary Assistant to the Academic Dean at Santa Barbara Business College and secretary at San Jose State University Police Department. From 1986 to 1989, she was a full-time student at San Jose State University and graduated in May 1990.

**WALTER F. TICHY, Visiting Scientist (1988)**

Tichy joined RIACS in August, 1988, from the University of Karlsruhe, where he is a Professor of Computer Science. Tichy received his B.S. from the Technical University in Munich in 1974 and his M.S. and Ph.D. degrees in Computer Science from Carnegie-Mellon University in 1976 and 1980, respectively. Previously, he was a Senior Scientist at Carnegie Group, Inc., and Associate Professor of Computer Science at Purdue University. His primary interest include large-scale parallelism, software engineering, software tools, and programming languages. He is a member of the ACM, the GI, and the IEEE and Sigma Xi.

**RAYMOND TUMINARO, Student (1986-89)**

Tuminaro came to RIACS in June, 1986, as a Ph.D. student at Stanford University, from where he graduated in 1989. He worked with Tony Chan on the application of multigrid methods to parallel processing systems. During summer 1986, he implemented a multigrid method on the Intel iPSC hypercube, demonstrating the potential performance improvements using that architecture. His central interests include the development of computer algorithms to solve computational fluid dynamics problems on parallel computers. He is presently working on his Ph.D. in

the Numerical Analysis program in Computer Science at Stanford. He received a B.S. in Computer Science from Cornell University in 1984. He has been a teaching assistant and research assistant at Stanford and worked as a numerical consultant at the Stanford Linear Accelerator.

**CARL WILLIAMS, Research Associate (1989)**

Williams joined RIACS in November, 1989. He received his B.S. in Applied Mathematics/Computer Science from Kent State University in 1989 and his M.S. degree in Computer Science in 1989. During the summer of 1989 he was a research consultant at the Institute for Computational Mathematics, Kent State. He worked on designing and developing different algorithms for solving systems of linear equations for exact solutions on the STARAN and various shared-memory multiprocessors. He was previously employed by Digital Equipment Corporation, designing an AI system using logic programming to correct circuit design errors, and Hewlett-Packard, where he worked on evaluating the performance of the Network File System. He is currently working with the Parallel Systems on various projects involving the use of the Intel iPSC/2 Hypercube machine. This will include the design and implementation of software for the Intel iPSC/2 Hypercube. His long term goals are in the area of numerical simulation of various CFD problems on the Hypercube.

**JAMES A. WOODS, Senior Research Associate (1990)**

Woods enrolled at RIACS in January, 1990, after twelve years on contract with NASA Ames Research Center. He received a B.A. (1974) and M.S. (1975) in computer science, both from the University of California at Berkeley. His master's thesis concerned orthogonal transform coding of music. Woods has made several contributions to standard Unix by [co]authoring the utilities *compress*, *fast e?grep*, and *fast find* – software which employs advances in algorithm design ranging from Ziv-Lempel coding to hybridized Boyer-Moore and regular expression search. Woods recently was declared a winner in the Seventh International Obfuscated C Contest with one of the first truly useful Trojan Horse programs. Woods' original work on the combinatorics of anagrams has been cited in Knuth's *Art of Computer Programming* (volume 4, forthcoming) as well as in *Scientific American*.

# **Appendices**





---

## Appendix 1:

### Publications by RIACS Staff

#### Books and Book Chapters

1. **Peter J. Denning** (editor), *Computers Under Attack: Intruders, Worms, and Viruses*, Addison-Wesley (1990), in press.
2. **Richard F. Haines**, "Windows: Their Performance and Functions in Confining Environments," Chapter 31 in *From Antarctica to Outer Space: Life in Isolation and Confinement* (A. A. Harrison *et al.*, eds.), Springer-Verlag, New York (1990).
3. **Richard F. Haines**, "A Breakdown in Simultaneous Information Processing," Chapter 17 in *Presbyopia Research: From Molecular Biology to Visual Adaption* (G. Obrecht and L. Stark, eds.), Plenum Publishing Corp., New York (1991), in press.
4. **Marjory Johnson**, "Performance Analysis of CCSDS Path Service," in *Protocols for High-Speed Networks* (H. Rudin and R. Williamson, eds.), North Holland (1989), pp. 69-82.
5. **Marjory Johnson**, "Performance Issues in Management of the Space Station Information System," in *Integrated Network Management, I* (B. Meandzija and J. Westcott, eds.), North Holland (1989), pp. 423-434.

#### Reviewed Journals

1. **Wray Buntine** "Modeling Default and Likelihood Reasoning as Probabilistic Reasoning," *Annals of Mathematics and AI*, 1990.
2. **Wray Buntine**, "Learning Trees," *AI and Statistics 3rd International Workshop*, submitted.
3. **Peter Cheeseman** and Bob Kanefsky, "Evolutionary Tree Reconstruction," *Proceedings of the American Association Artificial Intelligence (AAAI) Spring Symposium on Minimum Message Length Encoding*, Spring 1990, Stanford University.
4. **Peter Cheeseman**, J. Goebel, K. Volk, F. Gerbault, Matthew Self, John Stutz, and Will Taylor, "A Bayesian Classification of the IRAS LRS Atlas," *Astronomy and Astrophysics* 222, L5-L8 (1989).
5. **Peter Cheeseman**, "On The Importance of Evidence: A Response to Halpern," *Computational Intelligence* (1990), to appear.
6. **Peter J. Denning**, "Massive Parallelism in the Future of Science," *American Scientist* 77, 1 (January-February 1989), 16-18.
7. **Peter J. Denning**, "Computing as a Discipline" (with D. Comer, D. Gries, M. Mulder, A. Tucker, J. Turner, P. Young), *Communications of ACM* 32, 1 (January 1989), 9-23. A condensed version was published in *IEEE Computer*, February 1989.
8. **Peter J. Denning**, "The Internet Worm," *American Scientist* 77, 2 (March-April 1989), 126-128.
9. **Peter J. Denning**, "Bayesian Learning," *American Scientist* 77, 3 (May-June 1989), 216-218.
10. **Peter J. Denning**, "Sparse Distributed Memory," *American Scientist* 77, 4 (July-August 1989), 333-335.
11. **Peter J. Denning**, "Worldnet," *American Scientist* 77, 5 (September-October 1989), 432-434.
12. **Peter J. Denning**, "The ARPANET After Twenty Years," *American Scientist* 77, 6 (November-December 1989), 530-534.

13. **Peter J. Denning**, "Stopping Computer Crimes," *American Scientist* 78, 1 (January-February 1990), 10-12.
14. **Peter J. Denning**, "Is Thinking Computable?" *American Scientist* 78, 2 (March-April 1990), 100-102.
15. **Peter J. Denning**, "Changing the Guard," *American Scientist* 78, 3 (May-June 1990).
16. **Peter J. Denning**, "About Time," *American Scientist* 78, 4 (July-August 1990), to appear.
17. **Peter J. Denning**, "Saving All the Bits," *American Scientist* 78, 5 (September-October 1990), to appear.
18. **Peter J. Denning** and **Walter F. Tichy**, "Highly Parallel Computation," *Science* magazine, to appear during fall 1990.
19. **Roland Freund**, "On Polynomial Approximations to  $f_a(z) = (z-a)^{-1}$  With Complex  $a$  and Some Applications to Certain Non-Hermitian Matrices," *Approximation Theory and its Applications* 5, (1989), pp. 15-31.
20. **Roland Freund** and **Thomas Huckle**, "Extension Problems for H-Unitary Matrices and Hermitian Block Toeplitz Matrices," *Zeitschrift fuer Angewandte Mathematik und Mechanik* 69 (1989), pp. 210-212.
21. **Roland Freund** and **Thomas Huckle**, "On H-Contractions and the Extension Problem for Hermitian block Toeplitz matrices," *Linear and Multilinear Algebra* 25, (1989), pp. 27-37.
22. **Roland Freund**, "On Conjugate Gradient Type Methods and Polynomial Preconditioners for a Class of Complex Non-Hermitian Matrices," *Numerische Mathematik*, accepted.
23. **Roland Freund** and **Bernd Fischer**, "Chebyshev Polynomials Are Not Always Optimal," *Journal of Approximation Theory*, accepted.
24. **Roland Freund**, "On polynomial preconditioning and asymptotic convergence factors for indefinite Hermitian matrices," *Linear Algebra and its Applications*, accepted.
25. **Steve Hammond** and **Robert Schreiber**, "Efficient ICCG on a Shared Memory Multiprocessor," *International Journal of High Speed Computing*, accepted.
26. **Marjory Johnson**, "Coping with Data from Space Station Freedom," *Computer Networks and ISDN Systems*, to appear.
27. **David Rogers**, "Kanerva's Sparse Distributed Memory: An Associative Memory Algorithm Well-Suited To The Connection Machine," *International Journal of High Speed Computing* 1, 2 (1989), 349-365.
28. **Youcef Saad** and **Peter N. Brown**, "Hybrid Krylov Methods for Nonlinear Systems of Equations," *SIAM Journal on Statistical and Scientific Computing* 11, No. 3, 450-481.
29. **Youcef Saad** and **S. Gallopoulos**, "Complexity of Fast Elliptic Solvers on Parallel Computers," *International Journal of High Speed Computing* 1 (1989), pp. 113-142.
30. **Youcef Saad** and **E. Anderson**, "Solving Sparse Triangular Systems on Parallel Computers," *International Journal of High Speed Computing* 1 (1989), pp. 73-96.
31. **Youcef Saad**, "Numerical Solution of Large Nonsymmetric Eigenvalue Problems," *Computer Physics Communications* 53 (1989), pp. 71-90.
32. **Youcef Saad**, "Krylov Subspace Methods on Supercomputers," *SIAM Journal on Statistical and Scientific Computing* 10 (1989), pp. 1200-1232.
33. **Youcef Saad** and **E. Gallopoulos**, "Parallel Block Cyclic Reduction Algorithms," *Parallel Computing* 10, (April 1989), pp. 143-160.
34. **Youcef Saad** and **M.H. Schultz**, "Data Communication in Hypercubes," *J. Dist. & Par. Computing* 6 (1989), pp. 115-135.
35. **Youcef Saad** and **M.H. Schultz**, "Data Communication in Parallel Architectures," *Parallel Computing* 11 (1989), pp. 131-150.

36. **Robert Schreiber** and Charles Van Loan, "A Storage Efficient WY Representation for Products of Householder Transformations," *SIAM Journal on Scientific and Statistical Computing* 10, 1 (1989).
37. **Robert Schreiber** and Gautam Shroff, "On the Convergence of the Cyclic Jacobi Method for Parallel Block Orderings," *SIAM Journal on Matrix Analysis and Applications* 10, 3 (1989), pp. 326-346.
38. **Robert Schreiber** and Joshua Ronen, "Seismic Data Processing on an SIMD Array Parallel Supercomputer," *Supercomputers in Seismic Exploration, SEG Handbook of Geophysical Exploration* 21, Elmer Eisner, ed., Pergamon Press, Oxford (1989).
39. **Robert Schreiber**, "Bidiagonalization and Symmetric Tridiagonalization by Systolic Array," *Journal of VLSI Signal Processing* 1, No. 5 (1990), 279-285.
40. **Robert Schreiber** and Lloyd N. Trefethen, "Average Case Stability of Gaussian Elimination," *SIAM Journal on Matrix Analysis and Applications* 11, No. 3 (1990).
41. **Robert Schreiber** and Nicholas J. Higham, "Fast Polar Decomposition of an Arbitrary Matrix," *SIAM Journal on Scientific and Statistical Computing* 11, No. 4 (1990), 648-655.
42. **Robert S. Schreiber** and J. Dongarra, "Automatic blocking of nested loops," *J. Parallel and Distributed Computing*, submitted.
43. **Robert S. Schreiber** and J. Gilbert, "Optimal Expression Evaluation for Data Parallel Architectures," *J. Parallel and Distributed Computing*, submitted.
44. **Robert S. Schreiber** and V. Pan, "An Improved Newton Iteration for the Generalized Inverse of a Matrix, with Applications," *SIAM J. Scientific and Statistical Computing*, to appear.
45. **Robert S. Schreiber** and J. Gilbert, "Massively Parallel Sparse Cholesky Factorization," *SIAM J. Scientific and Statistical Computing*, submitted.
46. **Robert S. Schreiber** and F. Alvarado, "Optimal Parallel Solution of Triangular Sparse Systems," *SIAM J. Scientific and Statistical Computing*, submitted.
47. **Robert Schreiber** and V. Pan, "Fast, Preconditioned Conjugate Gradient Toeplitz Solver," *International J. High Speed Computing*, submitted.
48. **Paul N. Swartztrauber**, "Fast Fractional Fourier Transform and Applications," December 1989, *SIAM Journal*, submitted.

## Reviewed Conferences

1. **Wray Buntine**, "Myths and Legends in Learning Classification Rules," *Conference Proceedings of AAAI-90*, Morgan Kaufmann, Boston, MA; July-August, 1990, to appear.
2. **Peter Cheeseman**, Matthew Self, Jim Kelly, Will Taylor, Don Freeman, and John Stutz, "Bayesian Classification," *Proceedings of AAAI Seventh Annual Conference on Artificial Intelligence* 2, AAAI Press, pp. 607-611.
3. **Peter Cheeseman** and Bob Kanefsky, "Evolutionary Tree Reconstruction," *Proceedings of the AAAI Spring Symposium on Minimum Message Length Encoding* (Spring 1990), AAAI Press, pp. 35-39.
4. **Roland Freund**, "Explicitly Solvable Complex Approximation Problems Related to Sine Polynomials," *Proceedings of the Sixth Texas International Symposium on Approximation Theory*, C.K. Chui, L.L. Schumaker, and D.J. Ward (eds.), Boston: Academic Press (1989), pp. 263-266.
5. **Roland Freund** and Bernd Fischer, "Optimal Chebyshev Polynomials on Ellipses in the Complex Plane," *Proceedings of the Sixth Texas International Symposium on Approximation Theory*, C.K. Chui, L.L. Schumaker, and D.J. Ward (eds.), Boston: Academic Press (1989), pp. 251-254.
6. **Roland Freund**, "On Bernstein Type Inequalities and a Weighted Chebyshev Approximation Problem on Ellipses," *Proceedings of Computational Methods and Function Theory*, Valparaíso, Chile, March 1989. In St. Ruscheweyh, E. B. Saff, L. C. Salinas, and R. S. Varga (eds.), *Lecture Notes in Mathematics* 1435, Springer Verlag (1990), 45-55.

7. **Richard F. Haines**, "An Information Throughput Model for Complex, Transparent, Telescience Systems," *Advances in Human Factors/Ergonomics 12B*, Elsevier, New York (1989), 354-360.
8. **Marjory J. Johnson**, "Performance Issues in Management of the Space Station Information System," *The First International Symposium on Integrated Network Management*, Boston, MA, May 1989.
9. **Marjory J. Johnson**, "Performance Analysis of CCSDS Path Service," *IFIP Workshop on Protocols for High-Speed Networks*, Zurich, Switzerland, May 1989.
10. **Marjory J. Johnson**, "Space Station Freedom Operations Management System," *International Symposium on Local Communications Systems Management*, Canterbury, UK (September 1990), accepted.
11. **Pentti Kanerva**, "Contour-Map Encoding of Shape for Early Vision," to appear in D.S. Touretsky (ed.), *Advances in Neural Information Processing Systems 2*, San Mateo, California: Morgan Kaufmann, 1990.
12. **Pentti Kanerva**, "A Cerebellar-Model Associative Memory as a Generalized Random-Access Memory," in *Proceedings of IEEE COMPCON 89*, Session 21B, Cerebellar Models of Associative Memory, San Francisco, CA, February 27 - March 3, 1989, Washington, D.C.: IEEE Computer Society Press, (1989).
13. **Egon Loebner**, "Intelligent Network Management and Functional Cerebellum Synthesis," in *Proceedings of IEEE COMPCON 89*, Session 21B, Cerebellar Models of Associative Memory, San Francisco, CA, February 27 - March 3, 1989, Washington, D.C.: IEEE Computer Society Press, (1989). Work in collaboration with Sparse Distributed Memory group.
14. **David Rogers**, "Predicting Weather using a Genetic Memory: a Combination of Kanerva's Sparse Distributed Memory and Holland's Genetic Algorithms," in D.S. Touretsky (ed.), *Advances in Neural Information Processing Systems 2*, San Mateo, California: Morgan Kaufmann, (1990).
15. **David Rogers**, "Kanerva's Sparse Distributed Memory: An Associative Memory Algorithm Well-Suited to the Connection Machine," in *Proceedings of a Conference on Scientific Applications of the Connection Machine*, NASA Ames Research Center, CA, September 12-14, 1988: World Scientific Publishing (1989), 282-298.
16. **David Rogers**, "Statistical Prediction with Kanerva's Sparse Distributed Memory," in *Advances in Neural Information Processing Systems I*, San Mateo: Morgan Kaufmann (1989), 586-593.
17. **Yousef Saad** and E. Gallopoulos, "Efficient Parallel Solution of Parabolic Equations: Implicit Methods on the Cedar Multiclustur," *Proceedings ICS 90*. Also CSRD technical report 903, March 22, 1990.
18. **Roland Freund**, "On Conjugate Gradient Type Methods and Polynomial Preconditioners for a Class of Complex Non-Hermitian Matrices," *Numerische Mathematik 57*, 285-312, (1990).
19. **Roland Freund** and Bernd Fischer, "New Bernstein Type Inequalities for Polynomials on Ellipses," *Complex Variables*, accepted.
20. **Roland Freund**, "Explicitly Solvable Complex Approximation Problems Related to Sine Polynomials," *Proceedings of the Sixth Texas International Symposium on Approximation Theory*, C.K. Chui, L.L. Schumaker, and D.J. Ward (eds.), Boston: Academic Press (1989), 263-266.
21. **Roland Freund** and Bernd Fischer, "Optimal Chebyshev Polynomial on Ellipses in the Complex Plane," *Proceedings of the Sixth Texas International Symposium on Approximation Theory*, C.K. Chui, L.L. Schumaker, and D.J. Ward (eds.), Boston: Academic Press (1989,) 251-254.
22. **Yousef Saad** and H. Rajic, "Application of Krylov Subspace Methods in Fluid Dynamics," Nuclear Science and Engineering, (special issue), *Proceedings of the 1989 Conference on Advances in Nuclear Engineering Computation and Radiation Shielding*, April 1989, Santa Fe, NM. To appear.
23. **Yousef Saad** and E. Lee, M. Schultz, "A Three-Dimensional Wide Angle Wave Propagation Equation with Density Variations," *Second IMACS Symposium on Computational Acoustics*, Princeton, New Jersey, March 1989.

24. **Youcef Saad**, "Parallel Iterative Methods for Sparse Linear and Nonlinear Equations," *Proceedings for the 7th International Conference on Finite Elements Methods for Fluid Flow*, Huntsville, Alabama, April 1989.
25. **Youcef Saad**, "Numerical Solution of Large Lyapunov Equations," *Proceedings of the 1989 MTNS Conference*, Amsterdam, The Netherlands, June 1989.
26. **Youcef Saad**, "Overview of Krylov Subspace Methods with Applications to Control," *Proceedings of the 1989 MTNS Conference*, Amsterdam, The Netherlands, June 1989.
27. **Youcef Saad** and E. Gallopoulos, "On the Parallel Solution of Parabolic Equations," *Proceedings of the International Conference on Supercomputing*, De Groot ed., ACM press, 1989, Heraklion, Crete, June 1989.
28. **Youcef Saad**, "Numerical Solution of Large Lyapunov equations," *Proceedings of the 1989 MTNS Conference*, to appear.
29. **Youcef Saad**, "An Overview of Krylov Subspace Methods with Applications to Control," *Proceedings the 1989 MTNS conference*, to appear.
30. **Youcef Saad**, "Parallel Iterative Methods for Sparse Linear and Nonlinear Equations," *Proceedings of the 7th International Conference on Finite Elements Methods for Fluid Flow*, Huntsville, AL, April 3-7, 1989.
31. **Youcef Saad**, D. Lee, and M. Schultz, "A Three-Dimensional Wide Angle Wave Propagation Equation with Density Variations," *Second IMACS Symposium on Computational Acoustics*, Princeton, New Jersey, March 15-17 1989.
32. **Youcef Saad** and E. Gallopoulos, "On the Parallel Solution of Parabolic Equations," *Proceedings of the International Conference on Supercomputing*, Heraklion, Crete, June 6-10, 1989, De Groot ed., ACM press, 1989.
33. **Youcef Saad**, "Projection Methods for the Numerical Solution of Markov Chain Models," *Proceedings of First International Workshop on the Numerical Solution of Markov Chains*, Jan 8-10, 1990, Raleigh, North Carolina.
34. **Youcef Saad**, "Krylov Subspace Methods: Theory, Implementations and Applications," *9th International Conference on Computing Methods in Applied Sciences and Engineering*, Jan. 29 - Feb 2, 1990, Paris, France.
35. **Youcef Saad** and P. Brown, "Projection Methods for Solving Nonlinear Systems of Equations," *International Conference on Defects, Singularities and Patterns in Nematic Liquid Crystals*, in Orsay, France, May 28 - June 3, 1990.

## Other Conferences

1. **Richard F. Haines** and R. W. Jackson, "Video Image Compression and Remote Life Science Experiment Monitoring," *1990 Annual Aerospace Medical Association Meeting*, New Orleans, LA, May 13-17, 1990.
- ✓ 2. **Richard G. Johnson**, "A U.S. Strategy for Global Change Research," *AAS 27th Goddard Memorial Symposium on Global Environmental Change: The Role of Space in Understanding Earth*, Washington, D.C., March 1989.
3. **Barry M. Leiner** and W. Bostwick, "Toward a National Research Network," *4th International Conference on Supercomputing*, Santa Clara, CA, May 1989.
4. **Barry M. Leiner** and Daryl Rasmussen, "NASA Telescience Testbed Pilot Program," *40th IAF Congress*, Malaga, Spain, October 1989. [Also to appear in *Acta Astronautica*.]
5. **Barry M. Leiner**, **Robert L. Brown**, and **Richard F. Haines**, "Collaboration Technology and Space Science," *41st IAF Congress*, Berlin, Germany, October 1990, accepted. [Also to appear in *Acta Astronautica*.]

## Other Published Articles

1. **Peter Cheeseman, John Stutz, Matthew Self, William Taylor, J. Goebel, K. Volk, and H. Walker,** "Automatic Classification of Spectra from the Infrared Astronomical Satellite (IRAS)," NASA RP 1217, March 1989.
2. **Peter J. Denning,** "Whats Next? Partners in Thought," *ACM Communications* 32, 1, January 1989, pp. 7-8.
3. **Peter J. Denning,** "New Directions for the Communications," *ACM Communications* 32, 2, February 1989, pp. 164-165.
4. **Peter J. Denning,** "Human Error and the Search for Blame," *ACM Communications* 33, 1, January 1990, pp. 6-7.
5. **Robert S. Schreiber,** "Towards the Teraflop: Survey of Available Machines and Recommendations for 1990-1992. (Joint with H. Simon). NAS internal report. Preliminary draft completed.

## Technical Reports 1989

1. **Robert L. Brown, Dee Doyle, Richard F. Haines,, and Michael Slocum,** "Final Report of the Telescience Workstation Project," RIACS Technical Report TR 89.1, January 1989, 24 pp.
2. **David Rogers,** "Statistical Prediction with Kanerva's Sparse Distributed Memory," RIACS Technical Report TR 89.2, February 1989.
- ✓3. **Peter J. Denning,** "The Internet Worm," RIACS Report TR 89.3, January 1989, 12 pp.
- ✓4. **Roland Freund,** "Explicitly solvable complex Chebyshev approximation method related to sine polynomials," RIACS TR 89.4, February 1989.
5. **Roland Freund and Bernd Fischer,** "Optimal Chebyshev polynomials on ellipses in the complex plane," RIACS TR 89.5, February 1989.
- ✓6. **Robert L. Brown,** "A Distributed Program Composition System," RIACS Technical Report TR 89.6, February 1989, 121 pp.
7. **Barry M. Leiner,** "Telescience Testbed Pilot Program Final Report," Volume 1, RIACS Technical Report TR 89.7, March 1989.
8. **Barry M. Leiner,** "Telescience Testbed Pilot Program Final Report," Volume 2, RIACS Technical Report TR 89.8, March 1989.
9. **Barry M. Leiner,** "Telescience Testbed Pilot Program Final Report," Volume 3, RIACS Technical Report TR 89.9, March 1989.
10. **Paul O. Frederickson,** "Totally Parallel Multilevel Algorithms for Sparse Elliptic Systems," RIACS Technical Report TR 89.10, March 1989, 5 pp.
11. **Michael R. Raugh (Ed.),** *Cerebeller Models of Associative Memory: Three Papers from IEEE COMPCON Spring '89*, containing articles by James S. Albus, Pentti Kanerva, and Egon Loebner, (Reprinted, with permission, from *Proceedings of the 34th IEEE computer Society International Conference*, San Francisco, CA, Feb. 27 - Mar. 3, 1989 and from the *Physiological Society*, Oxford, England) RIACS Technical Report 89.11, March 1989.
- ✓12. **Peter J. Denning,** "Bayesian Learning," RIACS Technical Report TR 89.12 (March 1989), 11 pp.
- ✓13. **Marjory J. Johnson,** "Performance Analysis of CCSDS Path Service," RIACS Technical Report TR 89.13, April 1989 (revised July 1989), 14 pp.
- ✓14. **Robert Schreiber,** "Fast, Preconditioned Conjugate Gradient Toeplitz Solver," RIACS Technical Report TR 89.14, March 1989.
15. **Jack Dennis,** "The Paradigm Compiler: Mapping a Functional Language for the Connection Machine," RIACS Technical Report TR 89.15, April 1989, 15 pp.

16. E. Gallopoulos and Youcef Saad, "Some Fast Elliptic Solvers on Parallel Architectures and Their Complexities," RIACS Technical Report TR 89.16, April 1989.
- ✓ 17. Bernd Fischer and Roland Freund, "Chebyshev Polynomials Are Not Always Optimal," RIACS Technical Report TR 89.17, April 1989.
- ✓ 18. Henry A. Sowizral, "Distributed User Services for Supercomputers," RIACS Technical Report TR 89.18, April 1989, 19 pp.
- ✓ 19. E. Gallopoulos and Youcef Saad, "On the Parallel Solution of Parabolic Equations in the Complex Plane," RIACS Technical Report TR 89.19, May 1989.
- ✓ 20. Youcef Saad, "Numerical Solution of Large Lyapunov Equations," RIACS Technical Report TR 89.20, May 1989.
- ✓ 21. Roland Freund, "On the Bernstein Type Inequalities and a Weighted Chebyshev Approximation Problem on Ellipses," RIACS Technical Report TR 89.21, May 1989.
22. Peter J. Denning, "Sparse Distributed Memory," RIACS Technical Report TR 89.22, May 1989.
23. Youcef Saad, "Overview of Krylov Subspace Methods with Applications to Control Problems," RIACS Technical Report TR 89.23, May 1989.
24. Steven W. Hammond and Robert Schreiber, "Efficient ICCG on a Shared Memory Multiprocessor," RIACS Technical Report TR 89.24, May 1989.
25. Barry M. Leiner, "Policy Issues in Interconnecting Networks," RIACS Technical Report TR 89.25, June 1989, 44 pp. (also available as RFC 1124)
26. Peter J. Denning, "Worldnet," RIACS Technical Report TR 89.26, July 1989, 12 pp.
27. Umesh D. Joglekar, "Learning to Read Aloud: A Neural Network Approach Using Sparse Distributed Memory," RIACS Technical Report 89.27, July 1989.
28. Louis A. Jaeckel, "An Alternative Design for a Sparse Distributed Memory," RIACS Technical Report 89.28, July 1989.
29. Louis A. Jaeckel, "Some Methods of Encoding Simple Visual Images for Use with a Sparse Distributed Memory, with Applications to Character Recognition," RIACS Technical Report 89.29, July 1989.
30. Louis A. Jaeckel, "A Class of Designs for a Sparse Distributed Memory," RIACS Technical Report 89.30, July 1989.
31. Richard F. Haines, V. Johnson, K. Vogelsong and W. Froloff, "Ames Life Science Testbed Evaluation," RIACS Technical Report TR 89.31, July 1989, 33 pp.
32. Roland Freund, "On Polynomial Preconditioning for Indefinite Hermitian Matrices," RIACS Technical Report TR 89.32, August 1989, 33 pp.
33. Roland Freund, "Pseudo Ritz Values for Indefinite Hermitian Matrices" RIACS Technical Report TR 89.33, August 1989.
34. Walter F. Tichy and Christian G. Herter, "Modula-2: An Extension for Highly Parallel Processors," RIACS Technical Report TR 89.34, September 1989.
35. Gautam Schroff, "A Parallel Algorithm for Eigenvalues and Eigenvectors of a General Complex Matrix," RIACS Technical Report TR 89.35, June 1989.
36. C. C. Jay Kuo, Tony F. Chan, and Charles Tong, "Multilevel Filtering Elliptic Preconditioners," RIACS Technical Report TR 89.36.
37. Eugene Levin, "The Role of Graphics Superworkstations in a Supercomputing Environment," RIACS Technical Report TR 89.37, September 1989, 13 pp.
38. Peter J. Denning, "The ARPANET after Twenty Years," RIACS Technical Report TR 89.38, September 1989, 14 pp.
39. William J. Stewart, Bernard Philippe, and Youcef Saad, "Numerical Methods in Markov Chain Modeling," RIACS Technical Report TR 89.39, October 1989.

40. **Yousef Saad**, "Projection Methods for the Numerical Solution of Markov Chain Models," RIACS Technical Report TR 89.40, October 1989.
41. **Tony F. Chan and Ray S. Tuminaro**, "Analysis of a Parallel Multigrid Algorithm," RIACS Technical Report TR 89.41, October 1989, 23 pp.
42. **A. J. Wathen**, "Optimal Moving Grids for Time-dependent Partial Differential Equations," RIACS Technical Report TR 89.42, September 1989.
43. **Eugene Levin, Harry Partridge, and J. R. Stallcop**, "Collision Integrals and High Temperature Transport Properties for N-N, O-O, and N-O," RIACS Technical Report TR 89.43, November 1989, 30 pp.
44. **Leonardo Dagum**, "A Fast Sorting Algorithm for a Hypersonic Rarefied Flow Particle Simulation on the Connection Machine," RIACS Technical Report TR 89.44, November 1989.
45. **Barry M. Leiner**, "Smart Instruments and the National Collaboratory," RIACS Technical Report TR 89.45, November 1989, 12 pp.
46. **Peter J. Denning**, "Human Error and the Search for Blame," RIACS Technical Report TR 89.46, November 1989, 5 pp.
47. **Peter J. Denning**, "Stopping Computer Crimes," RIACS Technical Report TR 89.47, November 1989, 11 pp.
48. **Paul N. Swarztrauber and Richard K. Sato**, "Solving the Shallow Water Equations on the Cray XMP and the Connection Machine 2," RIACS Technical Report TR 89.48, December 1989.
49. **Paul N. Swarztrauber**, "An  $O(\log^2 N)$  Parallel Algorithm for Computing the Eigenvalues of a Symmetric Tridiagonal Matrix," RIACS Technical Report TR 89.49, December 1989.
50. **Charles Tong and Paul N. Swarztrauber**, "Ordered Fast Fourier Transforms on a Massively Parallel Hypercube Multiprocessor," RIACS Technical Report TR 89.50, December 1989.
51. **Paul N. Swarztrauber and David H. Bailey**, "Efficient Detection of a CW Signal with a Linear Frequency Drift," RIACS Technical Report TR 89.51, December 1989.
52. **Paul N. Swarztrauber**, "Fast Fractional Fourier Transform and Applications," RIACS Technical Report TR 89.52, December 1989.
53. **Michael J. Flynn, Pentti Kanerva, and Neil Bhadhamkar**, "Sparse Distributed Memory: Principles and Operation," RIACS Technical Report 89.53, December 1989. (Published concurrently at Stanford University's Computer Systems Laboratory as Technical Report CSL-TR89-400.)
54. **Roland Freund**, "Conjugate Gradient Type Methods for Linear Systems with Complex Symmetric Coefficient Matrices," RIACS Technical Report TR 89.54, December 1989, 27 pp.
55. **Roland Freund and Thomas Huckle**, "A Restricted Signature Normal Form for Hermitian Matrices, Quasi-spectral Decompositions, and Applications," RIACS Technical Report 89.55, December 1989, 22 pp.
56. Released as TR 90.16.
57. **Yousef Saad**, "Globally Convergent Techniques in Nonlinear Newton-Krylov Algorithms," RIACS Technical Report TR 89.57, November 1989, 37 pp.
58. **Yousef Saad**, "Krylov Subspace Methods: Theory, Implementations and Applications," RIACS Technical Report TR 89.58, December 1989, (To appear in *9th International Conference on Computing Methods in Applied Sciences and Engineering*, Jan 29 - Feb 2, 1990, Paris, France.)



## Technical Reports 1990

1. **Richard F. Haines**, "A System Performance Throughput Model Applicable to Advanced Manned Telescience Systems," RIACS Technical Report TR 90.1, January 1990, 22 pp.
2. **Peter J. Denning**, "Is Thinking Computable?" RIACS Report TR 90.2, January 1990, 15 pp.
3. **David Rogers**, "BIRD: A General Interface for Sparse Distributed Memory," RIACS Technical Report TR 90.3, January 1990, 102 pp.
4. **Bruno Olshausen and A. B. Watson**, "The *Cortex Transform* as an Image Processor for Sparse Distributed Memory: An Initial Study," RIACS Technical Report 90.4, February 1990.
5. **Pentti Kanerva**, "Contour-Map Encoding of Shape for Early Vision," RIACS Technical Report 90.5, February 1990.
6. **David Rogers**, "Weather Prediction using a Genetic Memory," RIACS Technical Report 90.6, February 1990.
7. **Youcef Saad and Harry Wijshoff**, "SPARK: A Benchmark Package for Sparse Computations," RIACS Technical Report TR 90.7, February 1990.
8. **Paul Frederickson and Tim Barth**, "Higher Order Solutions of the Euler Equations on Unstructured Grids Using Quadratic Reconstruction," RIACS Technical Report TR 90.8, February 1990.
9. **Jukka Vanhalla and Kimmo Kaski**, "Neural Networks and MIMD multiprocessors," RIACS Technical Report 90.9, March 1990.
10. **Richard F. Haines**, "Human Performance Measurement-Validation Procedures Applicable to Advanced Manned Telescience Systems," RIACS Technical Report TR 90.10, February 1990, 31 pp.
11. **Louis A. Jaekel**, "Recognition of Simple Visual Images Using a Sparse Distributed Memory: Some Implementations and Experiments," RIACS Technical Report 90.11, March 1990.
12. **Roland Freund and Bernd Fischer**, "New Bernstein Type Inequalities for Polynomials on Ellipses," RIACS Technical Report TR 90.12, March 1990, 21 pp.
13. **Youcef Saad**, "Analysis of Some Krylov Subspace Approximations to the Matrix Exponential Operator," RIACS Technical Report TR 90.13, March 1990.
14. **Youcef Saad and E. Gallopoulos**, "Efficient Solution of Parabolic Equations by Polynomial Approximation Methods," RIACS Technical Report TR 90.14, March 1990.
15. **Robert Schreiber and John Gilbert**, "Optimal Expression Evaluation for Data Parallel Architectures," RIACS Technical Report TR 90.15, March 1990.
16. **Robert Schreiber and V. Pan**, "An Improved Newton Iteration for the Generalized Inverse of a Matrix, with Applications," RIACS Technical Report TR 90.16, March 1990, 38 pp.
17. **Peter J. Denning**, "On Writing a Column," RIACS Report TR 90.17 (March 1990), 7 pp.
18. **Douglas Danforth**, "An Empirical Investigation of Sparse Distributed Memory using Discrete Speech Recognition," RIACS Technical Report TR 90.18, March 1990.
19. **Richard F. Haines and R. W. Jackson**, "Television Image Compression and Small Animal Remote Monitorings," RIACS Technical Report TR 90.19, April 1990, 17 pp.
20. **Youcef Saad**, "SparsKit: A Basic Tool Kit for Sparse Matrix Computations," RIACS Technical Report TR 90.20, May 1990.
21. **Paul Frederickson and Oliver McBryan**, "Normalized Convergence Rates for the PSMG Method," RIACS Technical Report TR 90.21, June 1990.
22. **Steve Hammond and Robert Schreiber**, "Solving Unstructured Grid Problems on Massively Parallel Computers," RIACS Technical Report TR 90.22, July 1990.
23. **Wray Buntine**, "Myths and Legends in Learning Classification Rules," RIACS Technical Report TR 90.23, May 1990.
24. **Robert L. Brown and Dee Doyle**, "An Application for Multi-person Task Synchronization," RIACS Technical Report TR 90.24, July 1990, 12 pp.

25. **Barry M. Leiner, Robert L. Brown, and Richard F. Haines**, "Collaboration Technology and Space Science," RIACS Technical Report TR 90.25, July 1990, 10 pp.
26. **Leonardo Dagum**, "On Suitability of the Connection Machine for Direct Particle Simulation," RIACS Technical Report TR 90.26, June 1990.
27. **Peter Cheeseman**, "Evolutionary Tree Reconstruction," RIACS Technical Report TR 90.27, March 1990.

## **Research Memoranda**

1. **Douglas G. Danforth**, "Speech Transcription Using Sparse Distributed Memory," RIACS Memorandum 89.2, 1989.
2. **Pentti Kanerva, Bruno Olshausen**, "Two-Dimensional Shape Recognition Using Sparse Distributed Memory," RIACS Memorandum 89.3, 1989.
3. **Egon Loebner, Coe Miles-Schlichting, and David Rogers**, "Engineering Aspects of the Cerebellum: A Collaboration," RIACS Memorandum 89.4, 1989.

---

## Appendix 2: Seminars, Colloquia, and Lectures

### RIACS/NAS Distinguished Lecture Series

1. Ken Kennedy (Rice University), "Compiling Fortran for the Intel Touchstone," March 1990.
2. Burton Smith (Tera Computer), "The Tera Supercomputer," April 1990.
3. Jack Dongarra (University of Tennessee), "SCHEDULE: Tools for Developing and Analyzing Parallel Fortran Programs," May 1990.
4. Justin Rattner (Intel Scientific Computer), "The Microprocessor's Destiny: Addressing Grand Challenges Beyond the Reach of Traditional Supercomputers," June 1990.

### Seminars Sponsored or Co-sponsored by RIACS

1. Timothy Barth (Computational Fluid Dynamics Branch, NASA Ames Research Center), "Finite Volume Schemes on Unstructured Meshes," January 3, 1989.
2. Nicholas J. Higham (Cornell University), "Matrix Nearness Problems and Applications," January 10, 1989.
3. Randy Bank (University of California, San Diego), "A Class of Iterative Methods for Solving Saddle Point Problems," January 31, 1989.
4. Ray Tuminaro (Stanford University), "Performance of a Parallel Code for the Euler Equations on Hypercube Computers," February 7, 1989.
5. Creon Levit (NAS Systems Division, NASA Ames Research Center), "Explicit and Implicit Solution of the Navier-Stokes Equations," February 9, 1989.
6. Michael Berry (Thinking Machines Corporation), "Fortran-8X: A Data Parallel Language," February 21, 1989.
7. Elizabeth Ong (University of Washington), "The 3D Linear Hierarchical Basis Preconditioner, It's Shared Memory Parallel Implementation," February 28, 1989.
8. Michael T. Heath (Oak Ridge National Laboratory), "Parallel Solution of Sparse Linear Systems," March 9, 1989.
9. Christian Bischof (Argonne National Laboratory), "Incremental Condition Estimation and Applications," March 9, 1989.
10. John R. Baumgardner (Los Alamos National Laboratory), "3-D Numerical Investigation of Mantle Dynamics with Circum-Pangea Subduction As An Initial Condition Using Multigrid Methods," March 14, 1989.
11. Richard F. Haines (RIACS), "Ames Life Science Telescience Testbed Study," March 21, 1989.
12. Jerry Yan (Information Services Division, NASA Ames Research Center), "Post-Game Analysis-Heuristic Resource Management on Concurrent Systems," March 28, 1989.
13. Robert Schreiber (RIACS), "Solving Linear Systems on the Connection Machine," April 4, 1989.

14. Rod Fatoohi (NAS Applied Research Office, NASA Ames Research Center), "Multitasking a 2-D Navier-Stokes Algorithm on the CRAY-2 and CRAY Y-MP," April 11, 1989.
15. Jack Dennis (MIT), "Mapping Programs for Data Parallel Execution on the Connection Machine," April 17, 1989.
16. Constantine D. Polychronopoulos (CSR, University of Illinois), "Auto-Scheduling Compilers: A Powerful Environment for Parallel Programming and Scheduling," April 18, 1989.
17. David Rogers (RIACS), "Kanerva's Sparse Distributed Memory: A Cerebellar Model of Associative Memory Well-Suited to the Connection Machine," April 25, 1989.
18. Leonardo Dagum (Department of Aeronautics and Astronautics, Stanford University and RIACS), "Implementation of Hypersonic Rarefied Flow Particle Simulation on the Connection Machine," May 2, 1989.
19. Gautam Shroff (Rensselaer Polytechnic Institute and RIACS), "Parallel Jacobi Methods for Non-Symmetric Eigenvalue Problems," May 9, 1989.
20. Lennart Johnsson (Thinking Machines Corporation), "Scientific Computations and Numerical Algorithms on the Connection Machine," May 10, 1989.
21. Steven W. Hammond (RIACS), "Efficient ICCG on a Shared Memory Multiprocessor," May 16, 1989.
22. Kenneth P. Birman (Cornell University), "The ISIS Toolkit: A Programming Environment for Building Distributed Software," May 24, 1989.
23. Randy Allen (Arden Computers), "Why Even Scalar Machines Need Vector Compilers," May 30, 1989.
24. Iain S. Duff (Harwell Laboratory and CERFACS), "Sparse Matrix Research at CERFACS," May 30, 1989.
25. Jerry Hunter (University of Arizona), "OASIS: A Display and Control Tool for Robots," June 1, 1989.
26. Gary Bedrosian (GE Corporate Research and Development), "Parallel Computing for RF Scattering Calculations," June 6, 1989.
27. Michael Mascagni (Mathematical Research Branch, N.I.D.D.K., The National Institute of Health), "Random Walks, Elliptic Equations, and Massively Parallel Computing," June 16, 1989.
28. Bruno Olshausen (RIACS), "Two-Dimensional Shape Recognition Using Sparse Distributed Memory: An Example of a Machine Vision Systems that Exploits Massive Parallelism for both High-Level and Low-level Processing," June 27, 1989.
29. Henry A. Sowizral (RIACS), "Accelerating the Execution of Discrete-Event Simulation by Using Multiple Processors and Time Warps," June 29, 1989.
30. Anoop Gupta (Stanford University), "Design of Large Scale Shared-Memory Multiprocessors," July 25, 1989.
31. David Keyes (Yale University), "Parallel Solution of Coupled Transport Equations through Domain Decomposition," July 26, 1989.
32. Lewis Hitchner (RIACS), "Image Synthesis in the Context of a Visual Model," July 25, 1989.
33. Richard F. Haines (RIACS), "Ames Life Science Telescience Testbed Experiment," July 27, 1989.
34. Tracy T. Bradley (Hughes Aircraft), "Requirements for FDDI Systems and Associated Performance Issues," July 28, 1989.
35. Walter F. Tichy (University of Karlsruhe, FRG and RIACS), "Programming Highly Parallel Computers," August 22, 1989.
36. Jacek Maitan, (Lockheed), "Intelligent Reconfigurable Distributed Systems," August 28, 1989.
37. Andy Wathen (University of Bristol, Stanford University and RIACS), "A Study of Optimal Grids for Time-Dependent Partial Differential Equations," August 29, 1989.

38. **Richard Johnson** (RIACS), "Global Environmental Change: Science and Policy Issues," October 13, 1989.
39. **Alan J. Laub** (University of California, Santa Barbara), "Riccati Equations and the Matrix Sign Function," October 17, 1989.
40. **Rod A. Fatoohi** (NASA Applied Research Office, NASA Ames Research Center), "Vector Performance Analysis of Three Supercomputers: Cray-2, Cray Y-MP and ETA-10Q," October 31, 1989.
41. **David H. Bailey** (NAS Applied Research Office, NASA Ames Research Center), "FFTs in External or Hierarchical Memory," October 31, 1989.
42. **Creon Levit** (NAS Applied Research Office, NASA Ames Research Center), "Parallel Solution of Pentadiagonal Systems Using Generalized Odd-Even Elimination," October 31, 1989.
43. **Tuomo Kauranne** (European Centre for Medium-Range Weather Forecasts, United Kingdom), "Asymptotic Parallelism of PDEs and Atmospheric Dynamics," November 2, 1989.
44. **John Greenstadt** (San Jose State University), "Cell Discretization of Nonselfadjoint Linear Elliptic P.D.E.'s," November 20, 1989.
45. **Harry Wijshoff** (CSR, University of Illinois and RIACS), "A Large-Grain Parallel Sparse System Solver," November 21, 1989.
46. **Gary Montry** (Myrias Computer Corporation), "The Myrias SPS-2 Parallel Processor and the Parallel Applications Management System," November 28, 1989.
47. **Binay Sugla** (Bell Laboratories), "CAPER: An Application Programming Environment for Multiprocessors," November 29, 1989.
48. **Ray Tuminaro** (Department of Computer Science, Stanford University and RIACS), "A Highly Parallel Multigrid-Like Method for the Solution of the Euler Equations," November 30, 1989.
49. **Chiping Li** (Naval Research Laboratory), "A Uniform Algorithm for Boundary and Interior Points for Computational Fluid Dynamics Applications on the Connection Machine," December 1, 1989.
50. **Paul N. Swarztrauber** (National Center for Atmospheric Research (NCAR) and RIACS), "An  $O(\log^2 N)$  Parallel Algorithm for Computing The Eigenvalues of a Symmetric Tridiagonal Matrix," December 5, 1989.
51. **G.F. Carey** (Aerospace Engineering and Engineering Mechanics, University of Texas), "Adaptive Grids and Vector Parallel Schemes," December 15, 1989.
52. **Richard Prager** (Engineering Department, Cambridge University), "Modified Kanerva Model: Results for Real-Time Word Recognition," January 8, 1990.
53. **Victor Eliashberg** (Universal Learning Systems, Inc.), "E-Machines: Universal Learning Neurocomputers with Context-Sensitive Associative Memory," January 10, 1990.
54. **Ren Breck**, (Notepad International, Inc.), "The Role of Advanced Telecommunications in Disaster Management," January 25, 1990.
55. **Denny Dahl** (Thinking Machines Corporation), "Mapping and Compiled Routing on the CM," February 1990.
56. **Jeff McDonald** (Thermosciences Division, Aerothermodynamics Branch of NASA Ames Research Center), "A Computationally Efficient Particle Simulation Method Suited to Vector Computer Architectures," February 1990.
57. **Robert B. Haber** (University of Illinois at Urbana-Champaign, National Center for Supercomputing Applications), "Distributed Systems for Real-Time, Interactive Visualization," February 1990.
58. **Kiat Chua** (Graduate Aeronautical Laboratory, California Institute of Technology), "Vortex Calculations in Two-and-Three Dimensions," February 1990.
59. **Kimmo Kaski** (University of Tampere, Finland), "Can the Capacity of SDM be increased, i.e. Importance of Redundancy in Neural Networks," February 9, 1990.

60. **Louis Jaeckel** (RIACS), "Learning Machines and Statistical Inference," February 21, 1990.
61. **Douglas G. Danforth** (RIACS), "Speech and Shape Recognition," March 8, 1990.
62. **Scott B. Baden** (Lawrence Berkeley Laboratory), "Programming Abstractions for Localized Scientific Calculations Running on Multiprocessors," March 1990.
63. **Charles Tong** (University of California, Los Angeles), "Parallel Preconditioned Conjugate Gradient Methods for Elliptic Problems and Their Parallel Implementations," March 1990.
64. **Al Globus** (Sterling Software, NASA Ames Research Center), "ISOLEV in the FAST Lane," March 22, 1990.
65. **Cleve Moler** (MathWorks), "MATLAB," March 1990.
66. **Tony J. Garrett** (School of Mathematical Sciences, University of Bath, United Kingdom), "Numerical Calculation of Eigenvalues of Large Unsymmetric Matrices Arising from Discretizations of Differential Equations," April 1990.
67. **Peter Shirley** (University of Illinois at Urbana-Champaign), "Visualization of Non-rectilinear Scalar Volumes," April 1990.
68. **Harry F. Jordan** (University of Colorado, Department of Electrical and Computer Engineering), "Time-Space Trade-Offs in Optical Computing," April 1990.
69. **Ken Jacobsen** (MasPar Computer Corporation in Sunnyvale, California), "A Data-Parallel Version of ARC3D for the MasPar MP-1," April 1990.
70. **Robert F. Lucas** (Supercomputing Research Center, Bowie, Maryland), "An Implementation of the Multifrontal Method on MIMD Vector Machines," April 1990.
71. **John M. Levesque** (Pacific-Sierra Research Corporation), "Programming Tools for the 90's," May 1990.
72. **Anthony T. Chronopoulos** (University of Minnesota), "On Squaring Krylov Subspace Iterative Methods for Nonsymmetric Linear Systems," May 1990.
73. **Geoffrey Chesshire** (IBM-TJ Watson Research Center, New York), "CMPGRD: A System for Grid Generation, PDE Discretization and Solution Using Composite Overlapping Grids," May 1990.
74. **Rod Fatoohi** (NAS Applied Research Office, NASA Ames Research Center), "Vector Performance Analysis of the NEC SX-2," May 1990.
75. **Bracy H. Elton** (University of California, Davis), "A Numerical Theory for Lattice Gas and Lattice Boltzmann Methods for Solving PDE's," June 1990.
76. **Robert C. Clauer**, "Developing Teleoperations Capability for the Sondre Stromfjord, Greenland Upper Atmospheric Research Facility," June 7, 1990.
77. **Timothy Mattson** (STRAND Software Technologies, Inc.), "STRAND: Parallel Programming Tool," June 1990.

---

## Appendix 3:

### Talks and Presentations

1. **Peter J. Denning**, "Massive Parallelism in the Future of Science," UCLA, February 28, 1989.
2. **Peter J. Denning**, "Computer Science Strategic Perspective," Computing Research Board Conference, October 11, 1989.
3. **Peter J. Denning**, "Worldnet," Purdue University, November 3, 1989.
4. **Peter J. Denning**, "Worldnet," University of California at Santa Barbara, November 30, 1989.
5. **Peter J. Denning**, "Worldnet," University of San Francisco, November 30, 1989.
6. **Peter J. Denning**, "New Designs for Computing Education," ACM Computer Science Conference, February 21, 1990.
7. **Peter J. Denning**, "Is Thinking Computable?" ACM Chapter Phoenix, March 21, 1990.
8. **Peter J. Denning**, "Worldnet," IEEE Communications Conference, Keynote, March 22, 1990.
9. **Peter J. Denning**, "Information Technologies for Astrophysics Circa 2000," NASA workshop, May 23, 1990.
10. **Paul Frederickson**, "Totally Parallel Multilevel Algorithms for Sparse Elliptic Systems," Fourth Conference on Hypercubes, Concurrent Computers and Applications, Monterey, California, March 1989.
11. **Paul Frederickson**, Horst Simon, "Totally Parallel Multilevel Algorithms," SIAM Symposium on Sparse Matrices, Salishan Lodge, Gleneden Beach, Oregon, May 1989.
12. **Roland Freund**, "Explicitly Solvable Complex Chebyshev Approximation Problems Related to Sine Polynomials," Sixth Texas Symposium on Approximation Theory, College Station, Texas, January 1989.
13. **Roland Freund**, Bernd Fischer (speaker), "Optimal Chebyshev Polynomials on Ellipses in the Complex Plane," Sixth Texas Symposium on Approximation Theory, College Station, Texas, January 1989.
14. **Roland Freund**, "Constrained Chebyshev Approximation Problems on Ellipses," Conference on Computational Methods and Function Theory, Valparaiso, Chile, invited speaker, March 1989.
15. **Roland Freund**, "Polynomial Preconditioners for Shifted Hermitian Matrices and Optimal Chebyshev Polynomials on Ellipses," Colloquium of the Department of Mathematics, University of California at San Diego, invited speaker, April 1989.
16. **Roland Freund**, "Polynomial Preconditioners for Hermitian and Certain Complex Non-Hermitian Matrices," Numerical Analysis Seminar, Department of Mathematics, University of California at Berkeley, invited speaker, April 1989.
17. **Roland Freund**, "Conjugate Gradient Type Methods for Linear Systems with Complex Coefficient Matrices," SIAM Symposium on Sparse Matrices, Salishan Lodge, Gleneden Beach, Oregon, invited speaker, May 1989.
18. **Roland Freund**, Youcef Saad, "A Comparison of Polynomial Based Iterative Methods for Large Non-Symmetric Linear Systems," SIAM Symposium on Sparse Matrices, Salishan Lodge, Gleneden Beach, Oregon, May 1989.
19. **Roland Freund**, "Polynomial Preconditioners for Hermitian and Certain Non-Hermitian Matrices," 1989 SIAM Annual Meeting, San Diego, California, July 1989.

20. **Roland Freund**, "Conjugate Gradient Type Methods for Linear Systems with Complex Symmetric Coefficient Matrices," Copper Mountain Conference on Iterative Methods, Copper Mountain, Colorado, invited speaker, April 1-5, 1990.
21. **Roland Freund**, "Use of the Complex Symmetric Lanczos Recursion for Solving Linear Systems," Midwest NA Day Conference, University of Illinois at Urbana-Champaign, invited speaker, April 1990.
22. **Roland Freund**, "Conjugate Gradient Type Methods for Two Families of Complex Non-Hermitian Matrices," seminar at the Center for Supercomputing Research and Development, University of Illinois at Urbana-Champaign, invited speaker, April 1990.
23. **Roland Freund**, "A Quasi-minimal Residual Method for Non-Hermitian Linear Systems," The Householder Symposium XI on Numerical Algebra, Tylosand, Sweden, invited speaker, June 18-22, 1990.
24. **Roland Freund**, "Conjugate Gradient Type Methods for Complex Symmetric Linear Systems," Colloquium at the Department of Mathematics, University of Wuerzburg, West Germany, invited speaker, June 1990.
25. **Steve Hammond**, "Solving Sparse Triangular Systems on a Shared Memory Multiprocessor," 5th Parallel Circus, RPI, Troy, New York, April 1989.
26. **Steve Hammond, Rob Schreiber**, "Efficient ICCG on a shared memory multiprocessor," SIAM Symposium on Sparse Matrices, Salishan Lodge, Gleneden Beach, Oregon, May 1989.
27. **Lew Hitchner**, "Visualization for Planetary Exploration," 2nd Annual NASA/Stanford Scientific Visualization Workshop, Stanford University, June 12, 1990.
28. **Marjory Johnson**, "Network Protocol Performance," invited tutorial, ACM SIGMETRICS and PERFORMANCE 89 Conference, May, 1989.
29. **Marjory Johnson**, "Network Protocol Performance," invited tutorial, ACM SIGMETRICS and PERFORMANCE 90 Conference, May, 1990.
30. **Michael R. Raugh**, "Three Applications of SDM: Two-Dimensional Shape Recognition, Discrete Speech Experiments, and Weather Prediction," Information Systems Laboratory seminar, Stanford University, April 26, 1990.
31. **Michael R. Raugh**, "A Mathematical Framework for Pattern Accuracy in Microlithography," Solid State Seminar AKA EE430, Department of Electrical Engineering, Stanford University, February 28, 1990.
32. **Michael R. Raugh**, "Two-dimensional Self-Calibration Developed for Ebeam Lithography," Vision Colloquium, Department of Psychology (jointly with Human Factors Division of NASA Ames Research Center), Stanford University, February 5, 1990.
33. **Michael R. Raugh**, "An Overview of Sparse Distributed Memory: The Project, Architecture, and Applications," MIND special interest group in neural networks, University of Texas, Dallas, Texas, April 21, 1990.
34. **Michael R. Raugh**, "Danforth's Discrete-Speech Experiments Using SDM," Speech and Image Understanding Labs, Texas Instruments Inc., Dallas, Texas, April 20, 1990.
35. **Michael R. Raugh**, "The Architecture of Sparse Distributed Memory," TI's special interest group in neural networks, Texas Instruments Inc., Dallas, Texas, April 19, 1990.
36. **Michael R. Raugh**, "Three Applications of SDM: Two-Dimensional Shape Recognition, Discrete-Speech Experiments, and Weather Prediction," Computer Research Laboratory Colloquium, Texas Instruments Inc., Dallas, Texas, April 19, 1990.
37. **Michael R. Raugh**, "The SDM Architecture and Weather Prediction," special two-hour presentation for Annual Review of Research, RICOH California Research Center, Menlo Park, February 16, 1990.
38. **Michael R. Raugh**, "The SDM Architecture and Weather Prediction," (including a presentation by David Rogers), For David Peterson and Code SGE, Ames Research Center, February 7, 1990.



- 2
39. **Michael R. Raugh**, "The Sparse Distributed Memory Project at RIACS," Information Technology Department, RICOH California Research Center, Menlo Park, January 23, 1990.
  40. **Yousef Saad**, "Preconditioners for Least Squares Problems," N.A. Seminar, University of California at Berkeley, March 1989.
  41. **Yousef Saad**, "Parallel Iterative Methods for Sparse Linear Systems," Math Seminar, University of California at Los Angeles, April 1989.
  42. **Yousef Saad**, "SparsKit: Basic Tool Kit for Sparse Matrix Computations," SIAM Symposium on Sparse Matrices, Salishan Lodge, Gleneden Beach, Oregon, May 1989.
  43. **Yousef Saad**, "Globally Converging Nonlinear Krylov Subspace Methods," INRIA Seminar, Paris, France, July 1989.
  44. **Yousef Saad**, Peter Brown (speaker), "Globally Convergent Krylov Subspace Methods," 1989 SIAM Annual Meeting, San Diego, California, July 1989.
  45. **Yousef Saad**, "Parallel Preconditioning for General Sparse Linear Systems," 1989 SIAM Annual Meeting, San Diego, California, July 1989.
  46. **Yousef Saad**, Robert Schreiber, "Norm Reducing Incomplete Factorization Techniques for General Sparse Linear Systems," 1989 SIAM Annual Meeting, San Diego, California, July 1989.
  47. **Yousef Saad**, Harry Wijshoff, "A Comparative Study of the Cray Family of Supercomputers Using a Sparse Benchmark Package," SIAM Symposium on Sparse Matrices, Salishan Lodge, Gleneden Beach, Oregon, May 1989.
  48. **Yousef Saad**, "Parallel Fast Poisson Solvers," Math Seminar, University of California at Los Angeles, October 1989.
  49. **Yousef Saad**, E. Gallopoulos, "Efficient Parallel Solution of Parabolic Equations: The Polynomial Approximation Approach," SIAM Conference on Parallel Processing, Chicago, Illinois, December, 1989.
  50. **Yousef Saad**, E. Gallopoulos (presenter), "Efficient Parallel Solution of Parabolic Equations: Implicit Methods," SIAM Conference on Parallel Processing, Chicago, Illinois, December, 1989.
  51. **Yousef Saad**, Harry Wijshoff (presenter), "Performance Study of Some Supercomputers Using a Sparse Matrix Benchmark," SIAM Conference on Parallel Processing, Chicago, Illinois, December, 1989.
  52. **Yousef Saad**, "Projection Methods for Solving Nonlinear Systems of Equations," International Conference on Defects, Singularities and Patterns in Nematic Liquid Crystals (NATO sponsored), Orsay, France May 28 through June 3, 1990.
  53. **Yousef Saad**, "Incomplete QR Factorizations and Other Preconditioners," Minisymposium on Iterative methods, University of Pittsburgh, invited speaker, April 7, 1990.
  54. **Yousef Saad**, "Incomplete QR Factorizations," "Standards for Sparse Matrix Software," and "Data Structures in SparsKit," three invited talks, Copper Mountain Conference on Iterative Methods, Copper Mountain, CO, April 2-5, 1990.
  55. **Yousef Saad**, "Krylov Subspace Methods: Theory, Algorithms and Applications," Ninth International Conference in Computing Methods in Applied Sciences and Engineering, INRIA, invited speaker, January 29 to February 2, 1990.
  56. **Yousef Saad**, "Krylov Subspace Methods: Theory, Algorithms and Applications," Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland, January 25-26, 1990.
  57. **Yousef Saad**, "Projection Methods for the Numerical Solution of Markov Chain Models," First International Conference on Numerical Solution of Markov Chains, North Carolina State University, invited speaker, January 8-10, 1990.
  58. **Yousef Saad**, E. Gallopoulos, "Efficient Parallel Solution of Parabolic Equations: The Polynomial Approximation Approach," SIAM Conference on Parallel Processing, Chicago, December 11-12, 1989.

59. **Yousef Saad, E. Gallopoulos (presenter), "Efficient Parallel Solution of Parabolic Equations: Implicit Methods, SIAM conference on Parallel Processing, Chicago, December 11-12, 1989.**
60. **Yousef Saad, R. Sincovec, P. Frederickson, H. Simon, R. Schreiber, "High Performance Matrix Algorithms tutorial," Supercomputing 89, November 13-17, 1989.**
61. **Yousef Saad, "Parallel Fast Poisson Solvers," UCLA Math Seminar, October 26, 1989.**
62. **Yousef Saad, P. Brown (speaker), "Globally Convergent Krylov Subspace Methods," SIAM National Meeting, San Diego, CA, July 17-21, 1989.**
63. **Yousef Saad, "Parallel Preconditionings for General Sparse Linear Systems," invited minisymposium presentation, SIAM National Meeting, San Diego, CA, July 17-21, 1989.**
64. **Yousef Saad, R. Schreiber, "Norm Reducing Incomplete Factorization Techniques for General Sparse Linear Systems," SIAM National Meeting, San Diego, CA, July 17-21, 1989.**
65. **Yousef Saad, two invited lectures at Trondheim Workshop on Numerical Linear Algebra June 12-18, 1989.**
66. **Yousef Saad, "Globally Converging Nonlinear Krylov Subspace Methods," INRIA, Paris, France, July 4, 1989.**
67. **Yousef Saad, R. Freund, "A Comparison of Polynomial Based Iterative Solvers," SIAM Symposium on Sparse Matrices, May 22-24, 1989.**
68. **Yousef Saad, "SparsKit: A Basic Tool Kit for Sparse Computations," SIAM Symposium on Sparse Matrices, May 22-24, 1989.**
69. **Yousef Saad, H. Wijshoff, "A Comparison Study of the Cray Family of Supercomputers Using a Sparse Matrix Benchmark," SIAM Symposium on Sparse Matrices, May 22-24, 1989.**
70. **Yousef Saad, "Parallel Iterative Methods for Sparse Linear Systems," UCLA Math Department, April 21, 1989.**
71. **Yousef Saad, "Preconditioners for Least Squares Problems," UC Berkeley NA Seminar, March 7, 1989.**
72. **Robert Schreiber, "Automatic Blocking for Memory Hierarchies," University of Washington, January 1989.**
73. **Robert Schreiber, "Parallel Algorithms in Signal Processing," SDI Sensor Signal Processing Workshop, Washington, D.C., April 1989.**
74. **Robert Schreiber, John Gilbert, "Highly Parallel Sparse Cholesky Factorization," SIAM Symposium on Sparse Matrices, Salishan Lodge, Gleneden Beach, Oregon, May 1989.**
75. **Robert Schreiber, "An Improved Newton Method for the Generalized Inverse of a Matrix With Applications," 1989 SIAM Annual Meeting, San Diego, California, July 1989.**
76. **Robert Schreiber, "LU Decomposition on the Connection Machine," 1989 SIAM Annual Meeting, San Diego, California, July 1989.**
77. **Robert Schreiber, "Highly Parallel Sparse Cholesky Factorization," US/Finnish Workshop on Advances in Scientific Computing, Helsinki, Finland, August 1989.**
78. **Robert Schreiber, "Highly Parallel Sparse Cholesky Factorization," Stanford University, Stanford, California, August 1989.**
79. **Robert S. Schreiber, "Theoretical Problems in Parallel Compiler Optimization," Caltech, January 1990.**
80. **Robert S. Schreiber, "Programming Data Parallel Machines," First Nordic Workshop on VLSI and Image Analysis, Linkoping, Sweden, March 8, 1990.**
81. **Robert S. Schreiber, "Programming Data Parallel Machines," University of Umea, Sweden, March 12, 1990.**
82. **Robert S. Schreiber, "Programming Data Parallel Machines," Computer Science Institute, University of Copenhagen, June 15, 1990.**

83. **Robert S. Schreiber**, "Linear Algebra and Massively Parallel Machines," University of Linköping, Sweden, March 9, 1990.
84. **Robert S. Schreiber**, "Linear Algebra and Massively Parallel Machines," University of Bergen, Norway, March 26, 1990.
85. **Robert S. Schreiber**, "Supercomputing in the 1990s," The Nordic Conference on Supercomputing, Oslo, March 1990.
86. **Richard Sincovec, Youcef Saad, Robert Schreiber, Paul Frederickson, Horst Simon**, "High Performance Matrix Algorithms for Large Scale Scientific and Engineering Computations," Supercomputing '89, Reno, Nevada, November, 1989.
87. **Richard Sincovec, Niel Madsen**, "The Independent Time Step Method," SIAM Conference on Parallel Processing, Chicago, Illinois, December, 1989.
88. **Richard F. Sincovec, Robert S. Schreiber**, "Parallel Computing in CFD," Tutorial at Parallel CFD 90, May 6-8, 1990.
89. **Richard F. Sincovec**, "The Independent Time Step Method," International Conference on Parallel Computing: Achievement, Problems, and Prospects, June 3-7, 1990.

---

## Appendix 4:

### Other Activities of RIACS Staff

1. **Peter J. Denning** served as Editor-in-Chief of the *ACM Communications*, the monthly flagship magazine.
2. **Peter J. Denning** served as chair of the ACM editorial committee.
3. **Peter J. Denning** received the distinguished service award from the Computing Research Association in February 1989.
4. **Peter J. Denning** received the distinguished service award from the Association for Computing Machinery (ACM) in February 1990.
5. **Richard F. Haines** was awarded U. S. Patent 4,932,098, "Grooming aid for collecting debris," June 12, 1990.
6. **Steve Hammond, Robert S. Schreiber,** and Gene Golub were coorganizers of 7th Parallel Circus, Stanford University, March 30-31, 1990.
7. **Marjory Johnson** continued as a member of the Consultative Committee for Space Data Systems.
8. **Marjory Johnson** is organizing the IFIP Workshop on Protocols for High Speed Networks to be held in November 1990.
9. **Marjory Johnson** was an official observer on the American National Standards Institute ANSC X3T9.5 Committee.
10. **Marjory Johnson** was host and organizer of the international meeting of Panel 1 of the Consultative Committee for Space Data Systems (CCSDS), April 1989.
11. **Barry Leiner** was a member of the Internet Activities Board.
12. **Barry Leiner** was a member of a workshop that produced a report to the NSF called "Toward a National Collaboratory," March 1989.
13. **Robert S. Schreiber** helped with the assignment of submitted papers to program committee members and refereeing some of the papers for the Supercomputing 90 conference in St. Louis.
14. **Robert S. Schreiber** participated in the ongoing MATLAB project and demonstrated a prototype at the Householder XI Meeting on Numerical Linear Algebra.
15. **Yousef Saad** organized a session on "Standards for Sparse Matrix Software" at the Copper Mountain Conference on Iterative Methods.
16. **Yousef Saad** co-organized a special session entitled: "Sparse Matrix Standards: BLAS2 and Beyond," at the SIAM conference on sparse matrices, May 22-24, 1989. Salishan, Oregon, 1989